

ARI Research Note 92-21

AD-A250 147



①

Knowledge Acquisition for Application: Cognitive Flexibility and Transfer of Training in Ill-Structured Domains

Rand J. Spiro, Paul J. Feltovich, Richard L. Coulson,
Michael Jacobson, Aydin Durgunoglu,
Susan Ravlin, and Jihn-Chang Jehng

University of Illinois

for

Contracting Officer's Representative
George W. Lawton

Office of Basic Research
Michael Kaplan, Director



April 1992



92-12416



United States Army
Research Institute for the Behavioral and Social Sciences

Approved for public release, distribution is unlimited

92 5 08 007

U.S. ARMY RESEARCH INSTITUTE FOR THE BEHAVIORAL AND SOCIAL SCIENCES

**A Field Operating Agency Under the Jurisdiction
of the Deputy Chief of Staff for Personnel**

EDGAR M. JOHNSON
Technical Director

MICHAEL D. SHALER
COL, AR
Commanding

Research accomplished under contract
for the Department of the Army

University of Illinois

Technical review by

George W. Lawton

NOTICES

DISTRIBUTION: This report has been cleared for release to the Defense Technical Information Center (DTIC) to comply with regulatory requirements. It has been given no primary distribution other than to DTIC and will be available only through DTIC or the National Technical Information Service (NTIS).

FINAL DISPOSITION: This report may be destroyed when it is no longer needed. Please do not return it to the U.S. Army Research Institute for the Behavioral and Social Sciences.

NOTE: The views, opinions, and findings in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other authorized documents.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
1a. REPORT SECURITY CLASSIFICATION Unclassified			1b. RESTRICTIVE MARKINGS ---		
2a. SECURITY CLASSIFICATION AUTHORITY ---			3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution is unlimited.		
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE ---					
4. PERFORMING ORGANIZATION REPORT NUMBER(S) ---			5. MONITORING ORGANIZATION REPORT NUMBER(S) ARI Research Note 92-21		
6a. NAME OF PERFORMING ORGANIZATION University of Illinois		6b. OFFICE SYMBOL (if applicable) ---		7a. NAME OF MONITORING ORGANIZATION U.S. Army Research Institute	
6c. ADDRESS (City, State, and ZIP Code) Center for the Study of Reading 51 Gerty Drive, Champaign, IL 61820			7b. ADDRESS (City, State, and ZIP Code) 5001 Eisenhower Avenue Alexandria, VA 22333-5600		
8a. NAME OF FUNDING/SPONSORING ORGANIZATION U.S. Army Research Institute for the Behavioral and Social Sciences		8b. OFFICE SYMBOL (if applicable) PERI-BR		9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER MDA903-86-K-0443	
9c. ADDRESS (City, State, and ZIP Code) 5001 Eisenhower Avenue Alexandria, VA 22333-5600			10. SOURCE OF FUNDING NUMBERS		
			PROGRAM ELEMENT NO. 61102B	PROJECT NO. 74F	TASK NO. N/A
			WORK UNIT ACCESSION NO. N/A		
11. TITLE (Include Security Classification) Knowledge Acquisition for Application: Cognitive Flexibility and Transfer of Training in Ill-Structured Domains					
12. PERSONAL AUTHOR(S) Spiro, Rand J.; Feltovich, Paul J.; Coulson, Richard L.; Jacobson, Michael; Durgunoglu, Aydin; Ravlin, Susan; and Jehng, Jihn-Chang					
13a. TYPE OF REPORT Final		13b. TIME COVERED FROM 86/10 TO 90/09		14. DATE OF REPORT (Year, Month, Day) 1992, April	
15. PAGE COUNT 240					
16. SUPPLEMENTARY NOTATION Contracting Officer's Representative: George Lawton					
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)		
FIELD	GROUP	SUB-GROUP	Learning Case-based learning		
			Cognitive flexibility Ill-structured domains		
			Transfer (Continued)		
19. ABSTRACT (Continue on reverse if necessary and identify by block number)					
<p>This report presents research findings concerned with advanced knowledge acquisition of difficult material in ill-structured domains. The program had three interdependent aspects: (1) patterns of learning failure in advanced treatments of complex material and their causes and consequences, (2) new theory-based remedies for the observed patterns of learning failure, and (3) hypertext instructional approaches for forestalling learning failure and developing cognitive flexibility based on findings from the first two parts of the program.</p> <p>In the first part of the report, research is presented that reveals a substantial incidence of learning failure that is the result of oversimplification. We illustrate some of the characteristics of advanced learning of complex conceptual material and note how these characteristics are often at odds with the goals and tactics of introductory instruction and with psychological biases in learning. We comprehensively catalog factors that contribute to suboptimal learning at the advanced stage and analyze the nature of resulting</p> <p>(Continued)</p>					
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION Unclassified		
22a. NAME OF RESPONSIBLE INDIVIDUAL George W. Lawton			22b. TELEPHONE (Include Area Code) (703) 274-5590		22c. OFFICE SYMBOL PERI-BR

ARI Research Note 92-21

18. SUBJECT TERMS (Continued)

Analogy	Videodisc
Hypermed	Application of knowledge
Misconception	Hypertext
Development	

19. ABSTRACT (Continued)

learning deficiencies. Special attention is devoted to findings concerning the way that local learning failures compound each other to build large-scale networks of durable and consequential misconception. Another important finding in this arena is that many learning failures seem to result from using strategies that are successful at earlier, introductory stages of learning, but not later; several findings are presented that indicate that what helps early in learning often hurts when more advanced understandings are attempted later. Detailed research on one kind of popular learning strategy, the use of analogy, illustrates the way early successes impede later learning of more advanced treatments of the same material.

In the second part of the report, which builds on the findings from the first part, we describe a theoretical orientation for more successful advanced knowledge acquisition in ill-structured domains--*Cognitive Flexibility Theory* (CFT). A learner who has developed "cognitive flexibility" will be able to use conceptual knowledge in an adaptive fashion, selecting and assembling disparate elements of knowledge to fit the needs of a particular understanding or problem-solving situation. CFT emphasizes the use of multiple mental and pedagogical representations (e.g., multiple analogies); the promotion of multiple alternative systems of linkage among knowledge elements; the promotion of schema assembly (as opposed to the retrieval of prepackaged schemas); the centrality of "cases of application" as a vehicle for engendering functional conceptual understanding; and the need for participatory learning, tutorial guidance, and adjunct support for aiding the management of complexity. Results indicating successful transfer when instruction follows the theory are presented.

The third part of the report introduces a computer hypertext approach that has been systematically developed based on the instructional implications of CFT. The theoretical underpinnings of the hypertext approach, used to produce *Cognitive Flexibility Hypertexts* (CFH) and its novel design features, are presented in detail and followed by a report of transfer gain from using a CFH. Finally, the CFH approach is illustrated in detail in a very complex and ill-structured domain, a hypertext for case-based learning of the structure of conceptual knowledge and its application in the military strategy of the "indirect approach."

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)



KNOWLEDGE ACQUISITION FOR APPLICATION: COGNITIVE FLEXIBILITY AND TRANSFER OF TRAINING IN ILL-STRUCTURED DOMAINS

CONTENTS

	Page
EXECUTIVE SUMMARY	6
Overview of Final Report	6
Summary of Major Accomplishments	8
Military Implications	10
COGNITIVE FLEXIBILITY THEORY FOR ADVANCED KNOWLEDGE ACQUISITION IN ILL-STRUCTURED DOMAINS: A DETAILED AND INTEGRATIVE SYNOPTIC OVERVIEW OF THE THEORETICAL ORIENTATION AND THE PROGRAM OF RESEARCH	12
Context: The Kinds of Knowledge Domains and Learning Purposes Studied	12
Deficiencies in Advanced Knowledge Acquisition	13
Cognitive Flexibility Theory: Themes of Advanced Knowledge Acquisition	15
Recapitulation: A Shift from Single to Multiple Repre- sentations and from Generic Schema Retrieval to Situation-Specific Knowledge Assembly	23
PATTERNS OF FAILURE IN ADVANCED KNOWLEDGE ACQUISITION IN COMPLEX DOMAINS	24
The Deep Structure of Complex Domains and the Development of Misconceptions: The Nature of Conceptual Understanding in Biomedicine	24
Analogy-Induced Misconception in Advanced Knowledge Acquisi- tion: A Detailed Study of One Kind of Maladaptive Reductive Bias	73
TOWARDS MORE SUCCESSFUL ADVANCED-STAGE LEARNING: COGNITIVE FLEXIBILITY THEORY AND INSTRUCTIONAL HYPERTEXT	86
Knowledge Acquisition for Application: Cognitive Flexibility and Transfer in Complex Content Domains	86
An Example of the Cognitive Flexibility Theory Principle of Multiple Representations: Multiple Analogies as Antidotes for Analogy-Driven Misconception	98

CONTENTS (Continued)

	Page
An Introductory Rationale for Hypertext: Cognitive Flexibility, Constructivism, and Random Access Instruction for Advanced Knowledge Acquisition in Ill-Structured Domains	116
Cognitive Flexibility, Random Access Instruction, and Hypertext: Theory and Technology for the Nonlinear and Multidimensional Traversal of Complex Subject Matter . .	129
Cognitive Flexibility, Hypertext, Computer-Assisted Instruction, and the Promotion of Transfer for Complex Knowledge	163
A Detailed Example of the Application of Cognitive Flexibility Theory to a Hypertext System for Case-Based Development of Conceptual Understanding in an Ill-Structured Domain: The Military Strategy of the Indirect Approach	170
ACKNOWLEDGMENTS	221
PAPERS AND PRESENTATIONS OF RESEARCH SUPPORTED BY THE GRANT . . .	223
REFERENCES	226

LIST OF FIGURES

Figure 1. Schematic representation of the sarcomere	40
2. Schematic representation of sarcomere during shortening	41
3. Quote from a second-year medical student	43
4. Quote from an established cardiovascular physician	44
5. Quote from a reasonably current textbook of clinical cardiology	45
6. Skeletal length-tension relationship	48
7. Medical student discussing the length-tension relationship	49

CONTENTS (Continued)

	Page
Figure 8. Why do students believe component 1? Converging influences	51
9. Medical student discussing the basis for the Frank-Starling cardiac function relationship	54
10. Cardiac length-tension relationship in comparison to skeletal length-tension relationship	56
11. Why do students believe component 2? Converging influences	57
12. Frank-Starling relationship	59
13. Medical student discussing the Frank-Starling relationship	60
14. Why do student believe component 3? Converging influences	62
15. Medical student discussing heart failure	64
16. Why do students believe component 4? Converging influences	66
17. The misconception; a network of reciprocating beliefs	68
18. Schematic representation of the sarcomere	100
19. Schematic representation of a sarcomere during shortening	101
20. Force production to create movement (contraction): The rowing crew analogy	102
21. The nature of overall gross movement of parts in relation to each other (pulling toward the middle): The analogy of two crews facing each other in a long boat	104

CONTENTS (Continued)

	Page
Figure 22. The nature of overall gross movement: The turnbuckle analogy	105
23. Structure and function of the casing (elastin) surrounding the force producers (i.e., sarco- meres are contained within the casing): The Chinese finger cuffs analogy	106
24. Recruitment of force producers: The galley ship analogy	107
25. The biochemical life process of the force producers	109
26. Scene 1	175
27. Scene 2	176
28. Scene 3	177
29. Scene 4	178
30. Scene 5	179
31. Scene 6	180
32. Scene 7	181
33. Scene 8	182
34. Scene 9	183
35. Background to scene 8	188
36. Themes for scene 8	190
37. Example of theme commentary	191
38. SIA theme-search menu	193
39. Map for scene 8	201
40. Combinatorial theme search	215

CONTENTS (Continued)

	Page
Figure 41. Chancellorsville scene	216
42. Chancellorsville scene: Background and description (part 1)	217
43. Chancellorsville scene: Background and description (part 2)	218

KNOWLEDGE ACQUISITION FOR APPLICATION: COGNITIVE FLEXIBILITY AND TRANSFER OF TRAINING IN ILL-STRUCTURED DOMAINS

1. Executive Summary

1.1 Overview of Final Report

This report presents the findings of a program of research concerned with advanced knowledge acquisition of difficult material in ill-structured domains. The research program had *three interdependent aspects*, which are presented in turn in the report:

- patterns of learning failure in advanced treatments of complex material and their causes and consequences;
- new theory-based remedies for the observed patterns of learning failure; and
- hypertext instructional approaches for forestalling learning failure and developing cognitive flexibility based on the findings from the first two parts of the program.

As is obvious from these three foci, the research program was very broad in its coverage of advanced learning, spanning in an integrated manner from *what goes wrong*, to a *detailed theory of how to fix what goes wrong* (or avoid its happening in the first place), to *highly specific instructional implications of the theory, implemented in innovative computer environments*. This breadth of coverage of issues related to advanced learning in complex domains allows the interlocking character of important issues to emerge in ways that are not possible when domains of inquiry are more restricted. For example, it helps greatly to know what is going wrong in order to know how to fix it; so a lot of research effort went into studying patterns of learning failure that later was rewarding in its implications for developing new approaches to achieving learning success. Each part of the program enhanced the other parts, permitting a cumulative story of advanced learning to emerge that is both unified and comprehensive. This Final Report tells that interwoven and encompassing story.

In the first part of the report (following a detailed synthesis of the overall program of study in Section 2, which should be read first), research is presented that has revealed a substantial incidence of learning failure attributable to various forms of oversimplification (Section 3). We illustrate some of the special characteristics of advanced learning of complex conceptual material, and we note how these characteristics are often at odds with the goals and tactics of introductory instruction and with psychological biases in learning. Factors that contribute to suboptimal learning at the advanced stage are comprehensively catalogued, and the nature of resulting learning deficiencies are analyzed in great detail. Special attention is devoted to findings concerning the way that local (and sometimes modest-appearing) learning failures compound each other to build large scale networks of durable and consequential misconception. Another important finding in this arena is that many learning failures appear to result from the use of strategies that are successful at earlier, more introductory stages of learning (where learning goals are less ambitious): several findings are presented that indicate that what helps early in learning often *hurts* when more advanced understandings are attempted later. Since most studies of learning strategies do not examine the later consequences of earlier learning, this important maladaptive pattern has gone largely undetected. Detailed research on one kind of popular learning strategy, the use of *analogy*, illustrates the way early successes may impede later learning of more advanced treatments of the same material. The findings from this first phase of the program are important for several reasons. The learning

patterns that were found to lead to failure were frequently unexpected; sometimes they represented approaches that are mistakenly assigned uniform *value* in learning and instruction. Identifying these pitfalls to advancing in the learning of a topic is important in avoiding them. You need to know what is going wrong in order to know how to fix it, and the second phase of the research program builds on the first one in just that way.

The second part of the research program is the subject of the first chapters of Section 4, where we describe a theoretical orientation for more successful advanced knowledge acquisition in ill-structured domains, Cognitive Flexibility Theory. *A learner who has developed "cognitive flexibility" will be able to utilize conceptual knowledge in an adaptive fashion, selecting and assembling disparate elements of knowledge to fit the needs of a particular understanding or problem solving situation (rather than relying on retrieving from memory knowledge that is monolithically prepackaged for use).* This is needed in ill-structured knowledge domains, where variability across cases in the elements of knowledge that must be used, and how that knowledge is used in a particular case, are too great to allow adequate performance using prepackaged scripts for thought and action that have to be applied to a wide range of cases -- ill-structuredness does not permit that kind of generalizability of abstract conceptual representations. The orientation of Cognitive Flexibility Theory emphasizes:

- the use of *multiple mental and pedagogical representations* (e.g., multiple analogies)-- the more representations or perspectives available to be drawn on, the fuller the coverage of individual cases during learning and the greater the flexibility in selecting a subset to best fit a new case encountered in the future;
- the promotion of *multiple alternative systems of linkage among knowledge elements*, providing flexibility in the routes that can be explored to assemble knowledge elements in different ways to best fit the case at hand;
- the *centrality of "cases of application"* as a vehicle for engendering functional conceptual understanding -- in an ill-structured domain, cases differ too much from each other, however they may be categorized, for any abstraction away from the individual case to be likely to capture enough of any single case's particulars -- conceptual knowledge must be taught in its context of application; and
- the need for *participatory learning*, tutorial guidance, and adjunct support for aiding the management of complexity.

Following Cognitive Flexibility Theory, learning proceeds by a nonlinear process of "*criss-crossing*" the cases in a knowledge domain in many conceptual directions, maximizing coverage within the domain (important because ill-structured domains give a different appearance in different of their sectors), increasing linkages across different parts of the domain, and instilling in the learner the importance of not being limited in the perspectives used to analyze the domain. This new approach to conceptualizing learning and instruction can be contrasted with the approaches of traditional texts and lectures, whose more linear and compartmentalized nature makes them more likely to miss the complexities and irregularities of ill-structured knowledge domains. Results indicating successful transfer for learners who are presented instructional material according to the tenets of Cognitive Flexibility Theory, compared to

control subjects receiving the same content presented with more traditional and conventional organizations, are presented.

The third part of the report, presented in the second half of Section 4, introduces a computer *hypertext* approach that has been systematically developed based on the instructional implications of Cognitive Flexibility Theory. The theoretical underpinnings and novel design features of the hypertext approach, used to produce *Cognitive Flexibility Hypertexts*, are presented in detail, followed by a report concerning transfer gain from using a Cognitive Flexibility Hypertext compared to learners in a control condition who see the same content presented in a more traditional computer learning environment lacking key features of Cognitive Flexibility Theory. Finally, the Cognitive Flexibility Hypertext approach receives a highly detailed illustration in an application of Cognitive Flexibility Theory to the analysis of instructional needs in a very complex and ill-structured domain, a hypertext for case-based learning of the structure of conceptual knowledge and its application in the *military strategy of the "indirect approach."* Although this strategy has been the basis of some of the most successful modern military campaigns and is thought by many contemporary expert military strategists to be the most useful one for the United States and its allies to pursue currently and in the coming decades (not to mention its prominence in U.S. strategy in the recent Persian Gulf War), the strategy is often discounted in military planning as unreliable because it is not directly teachable, as instead relating to the ineffable art or genius of command that is the lucky possession of some military decision makers. We argue that indeed SIA can not be taught--if it is taught conventionally as a well-structured domain. The last chapter shows how it can be taught in the very different way that characterizes Cognitive Flexibility Theory, using the new technological possibilities opened up by Cognitive Flexibility Hypertexts.

1.2 Summary of Major Accomplishments

The program of research presented in this Report has had many important successes in areas of learning and instructional theory and practice that have been in great need of progress. A full appreciation of these advances, which have been presented in major invited addresses to the American Psychological Association, Cognitive Science Society, American Educational Research Association, and other organizations (see Section 6), will of course require a careful reading of the Report. However, some of the more significant aspects of this work are highlighted below.

- A lot has been written about the *need* to better understand complexity and ill-structuredness, given the prevalence of such features in real-world knowledge domains. However, that acknowledged need has not been well addressed by systematic analysis of the features of complexity, and how those features relate to (especially how they make difficult) the attainment of advanced knowledge acquisition goals. Progress in this area has been a central accomplishment of this research effort.

- The cognitive characteristics of novices and experts have received ample study, but not much attention has been paid to the very important and lengthy stage of learning that comes between introductory learning and expertise. This neglected phase of learning received extensive study in our program. Our findings that early learning strategies and successes may seriously interfere with more advanced learning provides an important and previously barely

suspected indication of *barriers* to the later attainment of expertise (as well as indications of how those barriers can be avoided or surmounted). We now know that the path from novicehood to expertise is not one of monotonous progress, but rather requires a stage of *undoing* the results of earlier learning before further progress can be made.

- There has been some research on the existence of misconceptions in complex knowledge areas. However, insufficient attention has been devoted to such issues as: the different kinds of misconceptions and the ramifications of misconceptions of different types; the *sources* and *patterns of development* of misconceptions (e.g., in the *interaction* of learner biases, teaching strategies, testing practices, textbook writing styles, etc.), as well as the ways misconceptions relate to and influence *each other* in the ecology of knowledge domain representation. All of these important research questions received considerable illumination in our studies.

- A lot of attention has been devoted recently to statements about the need to embed learning in the kinds of situations that learning will have to be applied to. The present research falls within this developing movement of *situated learning*. However, our work goes beyond the state of the art in situated learning in two important ways. First, instead of markedly devaluing the importance of *explicit conceptual understanding* (perhaps on the assumption that it will either interfere with the extraction of situational information or that it will occur automatically from exposure to situations), we show how conceptual understanding can be systematically promoted while retaining the situatedness of its application. Second, much of the literature on situated learning says that we should do it, but very little is offered by way of articulating *specific programs for how to do it*. We have shown, with a great degree of detailed specificity, what programs of situated learning and instruction would actually look like. And we have provided some empirical support for the efficacy of this specific, theory-based approach in effecting improved transfer ability, historically a very difficult task in restricted laboratory studies (although we believe that our findings are just suggestive at this point, with much more study needed before we would argue that our claims have been definitively proven). The fact that the theory can receive the kind of specific and explicit instantiation as it does in our computer hypertext systems is an indication of its viability.

- *Hypertext* (nonlinear and multidimensional computer text) is one of the most frequently heard buzzwords in educational technology these days. However, instructional hypertext development has been too often characterized by *ad hoc* approaches that rely on the power of the technology and the ease of use of software like HyperCard to link instructional ideas in haphazard fashion with resulting difficulty of use by learners and dubious prospects for producing effective learning outcomes. Our hypertext development effort is different, and better, in several respects:

- First, our hypertext approach is *theory-based*--it is derived from studies of patterns of learning failure and a theory of learning and instruction tailored to the special needs of advanced knowledge acquisition in ill-structured domains.
- Second, the Cognitive Flexibility Hypertext approach to software design is *systematic* rather than arbitrary and accidental: the same very specific software design

recommendations are implemented in the same way in very different kinds of domains. This systematicity makes the approach much easier to *generalize* to new situations.

- Third, Cognitive Flexibility Hypertexts are *easier for learners to use* than other hypertexts. For example, the notorious hypertext problem of "getting lost in hyperspace" (lost in a labyrinth of connected pieces of information) is not a danger with Cognitive Flexibility Hypertexts.
- Fourth, Cognitive Flexibility Hypertexts should make design and development easier for programmers than other computer instructional systems (e.g., because of the specific guidelines, and because Cognitive Flexibility Hypertexts do not present a "scaling up" problem).
- Fifth, we have data indicating the Cognitive Flexibility Hypertext approach leads to greater transfer than more traditional computer-aided learning approaches that present the same content (although, again, we would not consider the empirical support definitive by any means--something as new as this approach is can only proceed one step at a time; however, the first, necessarily simplest steps have given no reason for discouragement and a basis for optimism about future successes).
- Sixth, the fact that we have been able to develop working prototype Cognitive Flexibility Hypertexts for real domains proves that it can be done, that the application of Cognitive Flexibility Theory to actual hypertext design is not just an impractical dream. Hypertext has the potential to revolutionize learning in complex knowledge areas. Our systematic, specific, and theory-based program of hypertext research and development should provide one of the first bases for building a *science* of hypertext--a solid foundation for further advance.

These were some of the more important accomplishments of our program of research. Again, however, the preceding list provides just the barest indication of the richness of our findings and their significance, and the interested reader is strongly encouraged to pursue topics of interest in their detailed presentations throughout this Report.

1.3 Military Implications

Not enough is known about the processes involved in moving a learner from the knowledge acquired in introductory-level classroom instruction to more advanced stages of learning, and then to the independent application (or transfer) of that knowledge to naturally occurring cases. The problem of knowledge application is especially great in complex and ill-structured domains, where concepts are harder to learn and where generic rules tend to be more restricted in their scope of application. Every kind of military training situation has some component that produces difficulty because of its complexity or ill-structuredness. Successful performance in these areas tends to come only with a considerable accumulation of actual case experience. Application of the training principles we are proposing has the potential to reduce the amount of time required to develop advanced skill levels in these difficult areas.

The development of means to help learners and trainees to master complex domains and foster cognitive flexibility is especially crucial for the *modern* military, both because of changing training and performance needs for the individual soldier and changing doctrine. The increasing complexity of the training goals in the military (for example, associated with the increasing complexity of the equipment that must be operated) is hardly news. And in many personnel specialties the nature of the assignment has always required an ability to deal with complexity and to manage data flexibly (e.g., battlefield command; military intelligence). With respect to doctrine, with the absence of the previous threat from the Soviet Union in Central Europe, the military must be prepared to respond flexibly to developing situations around the world, without having a standing force and fully developed battle plans already at the ready. For example, at a recent meeting of NATO military chiefs, Colin Powell joined the others in endorsing a proposal to "make their forces more *adaptable to changing situations*. "What we are aiming at is *maximum flexibility*," said Norwegian Gen. V. Eide, chairman of the alliance's military committee. [emphasis added]" (Associated Press, 1990). Also with respect to doctrine, the dominant strategy of the current era, given limitations of resources and uncertainty as to the characteristics of the enemy and the arena of conflict, is likely to be that employed in the recent Persian Gulf War, which can be described as a strategy of indirect approach. Not only does this strategy require commanders (in headquarters and in the battlefield) to flexibly adapt to dynamically changing situations in the way that Cognitive Flexibility Theory prepares for, but a military strategy of that kind is one of the domains that was studied in the present research.

Section 2. Cognitive Flexibility Theory for Advanced Knowledge Acquisition in Ill-Structured Domains: A Detailed and Integrative Synoptic Overview of the Theoretical Orientation and the Program of Research

We have indicated that we believe a major strength of this program of research is the high degree of integration and interconnectedness we have been able to attain across a very wide range of related areas of inquiry. The purpose of this section is to discuss the overall program in such a way as to be able to provide a readable *synoptic view* of its entire breadth that nevertheless presents sufficient detail to make the nature of each of the parts clear.

2.1 Context: The Kinds of Knowledge Domains and Learning Purposes Studied

One of our major theoretical claims, backed up by findings in our research, is that different approaches to learning and instruction will be required depending on characteristics of the knowledge domain that is the object of study and the purposes of learning. We have studied learning in *ill-structured domains* for purposes of *advanced knowledge acquisition*, two terms that will be taken up in turn.

2.1.1 Ill-Structuredness of Knowledge Domains. The concept of "ill-structuredness" is itself ill-structured. As will be made clear throughout the rest of this report, this does not mean that the term is so vague that it is useless in classifying domains. In fact, we can offer as a definition the following: by ill-structuredness we mean that many concepts (interacting contextually) are pertinent in the typical case of knowledge application, and that their patterns of combination are inconsistent across case applications of the same nominal type. However, we believe the term has additional meaning that emerges in the successive examples of how we use it, meaning that is in part determined by its context of use (as suggested by Wittgenstein, 1953). Thus, following the advice that will be found throughout this report, avoid oversimplification by reliance on abstractions when studying complex concepts that take on additional character when placed in context, the reader should hold the above definition of ill-structuredness as partial and provisional, pending further elaboration by exposures to its contexts of use. (See section 4.1.5.1 for a more detailed treatment of the nature and consequences of ill-structuredness.)

Physical and mathematical sciences tend to be more well-structured than biological science, which in turn tends to be more well-structured than social sciences. However, all domains contain substantial elements of ill-structuredness. Furthermore, and any domain of unconstrained real-world knowledge application will be ill-structured. So, mechanics or the rules of simple arithmetic or the techniques of statistics are well-structured; however, engineering to build a bridge using principles of mechanics, solving word problems using arithmetic or using statistics to analyze real data sets are all ill-structured situations.

2.1.2 The Goals of "Advanced Knowledge Acquisition". In our work we have been interested in "advanced knowledge acquisition"--learning beyond the introductory stage for a subject area, but before the achievement of practiced expertise that comes with massive experience. This often neglected intermediate stage is important because the aims and means of advanced knowledge acquisition are different from those of introductory learning. In

introductory learning the goal is often mere exposure to content and the establishment of a general orientation to a field; objectives of assessment are likewise confined to the simple effects of exposure (e.g., recognition and recall). At some point in learning about a knowledge domain the goal must change; at some point students must "get it right." This is the stage of advanced knowledge acquisition (Feltovich, Spiro, & Coulson, in press; Spiro, Feltovich, & Coulson, in press; Spiro, Vispoel, Schmitz, Samarapungavan, & Boerger, 1987): *The learner must attain a deeper understanding of content material, reason with it, and apply it flexibly in diverse contexts.* Obstacles to advanced knowledge acquisition include conceptual complexity and the increasing *ill-structuredness* that comes into play with more advanced approaches to a subject area.

The methods of education in introductory and advanced learning seem, in many ways, to be at odds. For example, compartmentalizing knowledge, presenting clear instances (and not the many pertinent exceptions), and employing reproductive memory criteria are often in conflict with the realities of advanced learning--knowledge which is intertwined and dependent, has significant context-dependent variations, and requires the ability to respond flexibly to "messy" application situations. These discrepancies in aims and tactics (along with many others that we have observed) raise the possibility that introductory learning, even when it is "successful," lays foundations in knowledge and in an approach to learning that interfere with advanced acquisition. As we have seen repeatedly demonstrated, that possibility is an actuality (Coulson, Feltovich, & Spiro, 1986; Feltovich, et al., in press; Spiro, et al., in press; Spiro, et al., 1987).

Much of our research was conducted with medical students. Medical school is an archetype of an advanced knowledge acquisition setting (Feltovich, et al., in press). Medical students have already had introductory exposure to many of the subject areas of biological science that they go on to study in medical school, but they are certainly not yet expert. Furthermore, the goals of medical education are clearly those of advanced knowledge acquisition. Important aspects of conceptual complexity must now be mastered (superficial familiarity with key concepts is no longer sufficient); and the ability to apply knowledge from formal instruction to real-world cases is certainly something that is expected of those studying to be physicians.

2.2 Deficiencies in Advanced Knowledge Acquisition

The first major part of our research program concerns patterns of learning failure in advanced knowledge acquisition. In this part of the program our focus of study has been on medical students' learning, understanding and application of important but difficult biomedical science concepts. This effort has revealed widely held systematic misconceptions among students, despite their having been exposed to appropriate information (Coulson, et al., 1986; Feltovich, et al., in press; Spiro, et al., in press; Spiro, et al., 1987). Stubborn misconception, notwithstanding usual classroom efforts at instruction, have been found for difficult concepts in other areas as well (e.g., physics: Champaign, Gunstone, & Klopfer, 1985; White, 1984).

The biomedical misconceptions that we have identified are of various kinds (Feltovich, et al., in press; Spiro, et al., in press). These include *contentive* errors, often involving

overgeneralization; for example, areas of subject matter are seen as being more similar than they really are. Errors attributable to dysfunctional biases in *mental representation* are also observed; for example, dynamic processes are often represented more statically. *Prefigurative "world views"* that underlie learners' understanding processes also cause problems; for example, the presupposition that the world works in such a way that "parts add up to wholes" leads students to decompose complex processes into components that are treated (mistakenly) as independent. Furthermore, at all these levels misconceptions interact in *reciprocally supportive* ways, and combine to yield higher order misconceptions (Coulson, et al., 1986; Feltovich, et al., in press). Failures of understanding *compound* themselves, building up durable chains of larger scale misconception.

2.2.1 Reductive biases: The pervasive role of oversimplification in the development of misconceptions. A predominant share of the misconceptions (and networks of misconception) that we have identified reflect one or another kind of *oversimplification* of complex material--associated with learners' earlier experiences with introductory learning, and even influenced by many experiences with advanced learning. Misconceptions of advanced material result both from interference from earlier, simplified treatments of that material and from a prevailing mode of approaching the learning process in general that fosters simplificational strategies and leaves learners without an appropriate cognitive repertoire for the processing of complexity (Feltovich, et al., in press; Spiro et al., in press; Spiro, et al., 1987).

We have termed the general tendency to reduce important aspects of complexity the *reductive bias*. Several forms of the bias have been identified, selected examples of which follow (see Coulson, et al., 1986; Feltovich, et al., in press; Spiro, et al., in press) for *examples of biomedical misconceptions* corresponding to the types of reductive bias listed):

2.2.1.1. Oversimplification of complex and irregular structure. Superficial similarities among related phenomena are treated as unifying characteristics. Interacting components are treated as independent. Incomplete conceptual accounts are presented (or accepted by the learner) as being comprehensive. Instances that are referred to as belonging to the same generic category are treated in a uniform manner despite their being highly diverse. The irregular is treated as regular, the nonroutine as routine, the disorderly as orderly, the continuous as discrete, the dynamic as static, the multidimensional as unidimensional. (This first reductive bias is the most general one, encompassing many of the specific ones listed below.)

2.2.1.2. Overreliance on a single basis for mental representation. A single, encompassing representational logic is applied to complex concepts and phenomena that are inadequately covered by that logic. For example: Understanding of a new concept is reduced to the features of a (partially) analogous concept. New, highly divergent examples are understood by exclusive reference to a single prototype. A single schema or theory is proffered and preferred, despite the fact that its coverage is significantly incomplete. Complexly multifaceted content has its understanding narrowed to just those aspects covered by a single organizational scheme. And so on.

3. Overreliance on "top down" processing. Understanding and decision making in knowledge application situations (i.e., cases) rely too exclusively on generic abstractions (i.e., concepts,

theories, etc.); detailed knowledge of case structure is not used enough (i.e., knowledge of "how cases go," as well as reasoning from specific case precedents).

4. Context-independent conceptual representation. The contexts in which a concept is relevant are treated as having overly uniform characteristics. This promotes the representation of conceptual knowledge in a manner too abstract for effective application (i.e., without sufficient regard for the specifics of application in context). Concepts are insufficiently tailored to their uses; concepts are not recognized as relevant when, in fact, they are; and concepts are mistakenly judged to be relevant in contexts where they are not.

5. Overreliance on precompiled knowledge structures. *Fixed* protocols or rigidly *prepackaged* schemas are presented to learners and used by them as recipes for what to do in new cases.

6. Rigid compartmentalization of knowledge components. Components of knowledge that are in fact interdependent are treated as being separable from each other. Learners develop mistaken beliefs in the independence of the components. Relatedly, where knowledge components do function independently, it may nevertheless be the case that conveying relationships between their conceptual structures would aid understanding; these connections are not drawn. When components are interrelated, there is a tendency to use just one linkage scheme, thereby underrepresenting the richness of interconnection in the system and promoting narrow, doctrinaire viewpoints (see the problem of single representations).

7. Passive transmission of knowledge. Knowledge is preemptively encoded under a scheme determined by external authority (e.g., a textbook) or a scheme which facilitates delivery and use. Knowledge is "handed" to the learner. The preemptive encoding is passively received by the learner, and useful benefits that result from personalized knowledge representations, derivable from active exploration and involvement in the subject area, do not develop. When active, participatory learning is encouraged, adequate support for the management of increased indeterminacy and cognitive load is not provided (e.g., mentor guidance, memory aids, etc.).

The next section will outline our theoretical approach to *remedying* the problems of advanced knowledge acquisition that are caused by these reductive biases.

2.3 Cognitive Flexibility Theory: Themes of Advanced Knowledge Acquisition

Where has our research on the *problems* of advanced knowledge acquisition led us? To an overall theoretical orientation that in many ways derives its fundamental themes from the specific nature of those learning problems, as the problems relate to the characteristics of ill-structured domains and the special goals of advanced knowledge acquisition (i.e., mastery of conceptual complexity and knowledge application/transfer).

In this section we provide a brief discussion of our most fundamental, theoretically motivated remedies for the problems of advanced knowledge acquisition. The following themes constitute different facets of what we call *cognitive flexibility* (Spiro, et al., 1987). The themes are, in a sense, *conditions* for developing mastery of complexity and knowledge transferability. Each of the headlined theoretical commitments has received some form of implementation, either in our experiments or in our theory-based computer hypertext

systems. The themes are discussed schematically and in the abstract. Detailed development of theoretical rationales and examples of our concrete instantiations of the themes can be found throughout this Report. The role of these themes in fostering transfer has been found in a few studies discussed herein (e.g., see sections xxx). However, this pattern of empirical support must be considered tentative at this time, limited as tests have been by the novelty and complexity of the approach and the need to proceed in a step-by-step manner. More experimental studies are clearly required. However, confidence in the conclusions of Cognitive Flexibility Theory is substantially bolstered from several converging directions that go beyond the positive data on transfer following Cognitive Flexibility Theory principles that has been found: positive confirmatory trends in several other studies; considerable empirical evidence from the studies of learning failure that point clearly to the tenets of Cognitive Flexibility Theory as compensating mechanisms and structural features for those patterns of failure; and the demonstrations of feasibility provided by the development of specific and detailed Cognitive Flexibility Theory models of the same general type in several disparate domains and implemented in working hypertext instruction prototypes.

2.3.1. Avoidance of Oversimplification and Overregularization. Because of the strong bias towards oversimplification that we have observed, it is clear that advanced knowledge acquisition must place a high premium on making salient those ways that knowledge is not as simple and orderly as it might first seem in introductory treatments. Where the problem is so often a presumption of simplicity and regularity, the remedy is to take special measures to *demonstrate complexities and irregularities*. It is important to lay bare the limitations of initial, first pass understandings, to highlight exceptions, to show how the superficially similar is dissimilar and how superficial unities are broken. Where conceptual error frequently occurs from atomistic decomposition of complexly interacting information, followed by misguided attempts at "additive" reassembly of the decomposed elements, the remedy is to take pains to *highlight component interactions*, to clearly demonstrate the intricate patterns of *conceptual combination*.

This is a very general theme, encompassing many of the others that follow in this list. Cognitive flexibility involves the selective use of knowledge to *adaptively fit* the needs of understanding and decision making in a particular situation; the potential for maximally adaptive *knowledge assembly* depends on having available as full a representation of complexity to draw upon as possible.

2.3.2. Multiple Representations. Single representations (e.g., a single schema, organizational logic, line of argument, prototype, analogy, etc.) will miss important facets of complex concepts. Cognitive flexibility is dependent upon having a diversified repertoire of ways of thinking about a conceptual topic. Knowledge that will have to be used in many ways has to be learned, represented, and tried out (in application) in many ways.

The use of multiple representations is important at different levels. For example, we have found multiple analogies to be very useful in understanding complex individual concepts (Spiro, et al., in press; see the example below of force production by muscle fibers; see also Collins & Gentner, in press; White & Frederiksen, 1987). However, the importance of multiple representations may be even more important for larger units of analysis. For example, we have found that students' understandings of the entire domain of biomedical knowledge is adversely

)

affected by the tendency to use just one way of modeling the various phenomena they encounter, one that comes from the *metaphor of the machine*. This one "lens" leads them to take for granted certain issues related to the nature of explanations, the structure of mental models of functional systems, and so on. These students develop understandings that do not capture important aspects of the biomedical domain (e.g., inherently organic processes). Their understandings would be more complete if they were to augment the selective view that results from their mechanistic bias with other understandings that selectively emerge from the unique aspects of other cognitive "lenses," for example, from *organicist* metaphors (Feltovich, et al., in press).

The need for multiple representations applies not only to complex concepts, but to cases as well. In an ill-structured domain, cases (examples, occurrences, events--occasions of use of conceptual knowledge) tend to be complex and highly variable, one to the next. The complexity of cases requires that they be represented from multiple theoretical/conceptual perspectives--if cases are treated narrowly by characterizing them using a too limited subset of their relevant perspectives, the ability to process future cases will be limited. First, there will be an assumption that cases are simpler than they in fact are, and attempts to deal with new cases will prematurely conclude after they are only *partially* analyzed. Second, there will be insufficient preparedness to deal with the specific patterns of interaction of theoretical/conceptual perspectives within cases. Third, to the extent that performance in future cases will require reasoning from sets of precedent cases (which is always a greater need in ill-structured domains), the likelihood of having case representations available in prior knowledge which are maximally apt in their relation to some new case is lessened to the extent that cases are narrowly represented in memory. This is especially so when there is substantial across-case dissimilarity; the relative novelty of a new case in an ill-structured domain will require more elaborate efforts to find appropriate precedents--the wider the variety that is available, the better the chances of finding a fit.

2.3.2.1 An Example of Multiple Representations: Integrated Multiple Analogies for Complex Concepts. As we have said, our studies of medical students have indicated that one of the most serious contributors to the problems of advanced knowledge acquisition is the use of a single knowledge representation. Complex concepts can rarely be adequately represented using a single schema, theoretical perspective, line of exposition, and so on. Nevertheless, in practice, complex concepts frequently are represented in some single fashion, with substantial consequences.

Our remedy has been to approach learning in all of the domains that we have studied with the goal of promoting *multiple representations* (e.g., multiple precedent cases for a new case; multiple organizational schemes for representing the same content material in our computer hypertexts; etc.). Here we will briefly consider just the case of analogy. We have discovered a large number of misconceptions that result from the overextended application of analogies (Spiro, et al., in press). To combat the negative effects of a powerful and seductive single analogy, we employ *sets of integrated multiple analogies*. Whenever a source concept in an analogy is missing important aspects of a target concept, or the source concept is in some way misleading about the target concept, we introduce *another* analogy to counteract those specific negative effects of the earlier analogy.

So, where we find that misconceptions about the nature of force production by muscle fibers often develop because of a common analogy to the operation of rowing crews (sarcomere "arms" and "oars" both generate force by a kind of "pulling"), other analogies are introduced to mitigate the negative effects of the limited rowing crew analogy (Spiro, et al., in press). An analogy to turnbuckles corrects misleading notions about the nature of relative movement and the gross structures within the muscle. An analogy to "finger handcuffs" covers important information missing in the rowing crew analogy about limits of fiber length (the elastin covering on muscle fiber bundles constricts at long lengths, stopping extension in a manner similar to the cross-hatched finger cuffs when you try to pull a finger out of each end). And so on. A composite imaging technique that helps the user to integrate the multiple analogies, so that the correct aspects of each analogy can be selectively instantiated in relevant contexts of use of the target concept, has also been developed. The procedure facilitates the learning of a concept (through the pedagogical benefits of analogy), while maintaining the integrity of the concept's complexities (by using multiple analogies to cover the concept's multifacetedness and to vitiate the force of incorrect aspects of any single analogy). (Also see Bunstein, 1983.)

2.3.2.2. Theory-based hypertext systems to implement the themes of advanced knowledge acquisition in ill-structured domains: The importance of revisiting and rearranging in the development of multiple representations. Much of the work on computer hypertext systems has been driven by the power of the technology, rather than by a coherent view of the cognitive psychology of nonlinear and multidimensional learning and instruction. In contrast, our hypertext approaches have a basis in cognitive theory--they derive from the themes of Cognitive Flexibility Theory. And their realm of operation is specified; they are especially targeted at advanced knowledge acquisition in ill-structured domains. (There is no point in imposing the extra cognitive load of nonlinearity and multidimensionality if the domain being studied is simple and well-structured, or if the goals of learning are the more easily attainable ones of introductory treatments.) We will briefly characterize our approach to implementing Cognitive Flexibility Theory in computer hypertext systems.

Our hypertext systems build multiple representations in a manner that can be understood using a metaphor of landscape exploration. Deep understanding of a complex landscape will not be obtained from a single traversal. Similarly for a *conceptual* landscape. Rather, *the landscape must be criss-crossed in many directions* to master its complexity and to avoid having the fullness of the domain attenuated (Spiro, et al., 1987; Wittgenstein, 1953). The same sites in a landscape (the same cases or concepts in a knowledge domain) should be *revisited* from different directions, thought about from different perspectives, and so on. There is a limit to how much understanding of a complex entity can be achieved in a single treatment, in a single context, for a single purpose. By *repeating* the presentation of the same complex case or concept information in *new contexts*, additional aspects of the multifacetedness of these "landscape sites" are brought out, enabling the kind of rich representations necessary in a complex and ill-structured domain. Thus, cognitive flexibility is fostered by a flexible approach to learning and instruction. The same content material is covered in different ways, at different times, in order to demonstrate the potential flexibility of use inherent in that content (Spiro & Jehng, in press; Spiro et al., 1987).

2.3.3. Centrality of Cases. The more ill-structured the domain, the poorer the guidance for knowledge *application* that "top-down" structures will generally provide. That is, the way

abstract concepts (theories, general principles, etc.) should be used to facilitate understanding and to dictate action in naturally occurring cases becomes increasingly indeterminate in ill-structured domains. The application of knowledge to cases in an ill-structured domain (i.e., a domain in which cases are individually multidimensional, and irregularly related one to the next) cannot be prescribed in advance by general principles. This is because, in ill-structured domains, there is great variability from case to case regarding which conceptual elements will be relevant and in what pattern of combination. In an ill-structured domain, general principles will not capture enough of the structured dynamics of cases; increased flexibility in responding to highly diverse new cases comes increasingly from reliance on reasoning from precedent cases.

Thus, examples/cases cannot be assigned the ancillary status of merely illustrating abstract principles (and then being discardable); the cases are key--examples are necessary, and not just nice (Feltovich, et al., in press; Spiro & Jehng, in press; Spiro, et al., 1987).

2.3.4. Conceptual Knowledge as Knowledge-in-Use. Not only is it more difficult to count on top down prescriptions for performance in new cases in an ill-structured domain (i.e., abstract concepts/theories inadequately determine responses to new cases), but there is also considerable indeterminateness in defining conditions for *accessing* conceptual structures in the first place, to engage the guidance the conceptual structures *do* offer. It is not that abstract knowledge has no role in ill-structured domains, but that its role is highly intertwined with that of case-centered reasoning. Put another way, in an ill-structured domain there will be greatly increased variability across cases in the way the same concept is used or applied. Thus it is harder to get from features of cases to the concepts that might need to be applied to those cases. And it is harder to apply a concept, once accessed, if it has many different kinds of uses across cases--concepts must be *tailored* to their application contexts. The Wittgensteinian dictum that meaning is determined by use clearly applies in ill-structured domains. If a concept's meaning in use cannot be determined universally across cases (as in an ill-structured domain), then one must pay much more attention to the details of how the concept is used--*knowledge in practice*, rather than in the abstract (Spiro & Jehng, in press; Spiro, et al., 1987; Wittgenstein, 1953).

In medical training, this issue of *variability and combination in concept instantiation* has an obvious implication for the traditional difficulty of integrating the biomedical basic science parts of the curriculum with the clinical parts. Physicians' practice would be improved if in problematic situations they could apply the interacting basic biomedical science concepts that underlie the clinical situation that is posing the problem. However, it is very difficult for medical students to learn how to get to the basic science concepts from clinical presenting features, partly because of the great variability across clinical cases in the way those concepts get instantiated. The *CardioWorld Explorer* hypertext system is designed to permit the learner to selectively examine the full range of uses of any selected basic science concept (or any selected combination of concepts) across cases with differing clinical features, teaching the patterns of concept application and thus facilitating *access* to conceptual information in clinical contexts (as well as fostering an understanding of the different ways that a given concept has to be tailored to be clinically relevant). [Our discussion of hypertext in this chapter focuses on the *CardioWorld Explorer* because the context of discussion primarily is using biological examples. Elsewhere in this report (see especially Chapters 4.3 - 4.6), other hypertext

systems derived from Cognitive Flexibility Theory that we have done more work on receive the bulk of attention, including one for the military strategy of the indirect approach.]

Again, in an ill-structured domain the meaning of a concept is intimately connected to its patterns of use. When the uses (instances, cases) of the same concept have a complex and irregular distribution (i.e., the domain is ill-structured), adequate prepackaged prescriptions for proper activation of the concept cannot be provided (i.e., *concept instantiation is non-routine*). Instead, greater weight (than in a well-structured domain) must be given to activating concepts in a new case by examination of family resemblances across the features of past cases that have been called (labeled as instances of) that concept.

2.3.5. Schema Assembly (from Rigidity to Flexibility). In an ill-structured domain, emphasis must be shifted from *retrieval* of intact, rigid, precompiled knowledge structures, to *assembly* of knowledge from different conceptual and precedent case sources to adaptively fit the situation at hand (Spiro, 1980; Spiro, et al., 1987). This follows, again, from characteristics of ill-structured domains. Since ill-structuredness implies kinds of complexity and irregularity that militate against the use of knowledge structures that assume routinizability across cases, the role of intact schema retrieval must be diminished--greater across-case differences cause a necessary decline in the ability of any large, single precompilation to fit a wide variety of cases. In complex and ill-structured domains, one cannot have a prepackaged schema for everything! As ill-structuredness increases, the use of *rigid* knowledge structures (i.e., the same precompiled knowledge structure used for many cases) must be replaced by *flexible*, recombinable knowledge structures. For any particular case, many *small* precompiled knowledge structures will need to be used. And there will be relatively little repetition of patterns across case-specific assemblies of these smaller pieces of precompiled knowledge. Accordingly, in knowledge acquisition for cognitive flexibility, the "storage of fixed knowledge is devalued in favor of the *mobilization of potential knowledge*" (Spiro, et al., 1987; see also Schank, 1982.)

2.3.6. Noncompartmentalization of Concepts and Cases (Multiple Interconnectedness). Because of the complex and irregular way that abstract conceptual features weave through cases/examples in ill-structured domains, knowledge cannot be neatly compartmentalized. In order to enable the situation-dependent, adaptive schema assembly from disparate knowledge sources that characterizes cognitive flexibility, those multiple sources must be highly interconnected. Concepts cannot be treated as separate "chapters." Retroactive assembly of independently taught, and noninterrelated, constituent conceptual aspects too often fails. Also, although cases have to be focused on separately, so that the complexity of case structure is conveyed, they should not be taught in *just* that way--connections across cases must also be established. Rather than relegating concepts or cases to separate compartments, chapters, and so on, our systems strive for *multiple interconnectedness (of cases and concepts) along multiple conceptual and clinical dimensions*.

Our approach to fostering multiple interconnectedness of knowledge representations in our hypertexts is to code case segments with a multidimensional vector indicating the relevance of a variety of thematic/conceptual dimensions to that case segment (Spiro & Jehng, in press). (Positive values in the vector also point to commentary, providing expert guidance about the nature of the conceptual dimension's instantiation in that particular case segment; this helps

with the problem of teaching conceptual knowledge-in-use discussed earlier). Then, as the hypertext program guides the learner in criss-crossing the domain's "landscape," by exploring patterns of overlap in the vectors for different case segments, knowledge representations are built up in which parts of cases are connected with many parts of other cases, along many conceptual/theoretical dimensions of case-segment similarity. In that way, many alternative paths are established to get from one part of the overall knowledge base to any other part of the knowledge base that aspects of some future case may signal as relevant. Thus, the potential for flexible, situation-adaptive schema assembly is fostered (along with such other virtues as the establishment of multiple routes for memory access to any node in the system).

So, for example, the *Cardioworld Explorer* hypertext is designed so that segments of clinical cases can be encoded with a vector of clinical and basic biomedical science themes that are relevant to each segment. The system can then establish connections between a segment of one case and segments of many other cases, along the various (conceptual and clinical) thematic dimensions represented in the vector. In case-based instruction, it is often true that there are important, instructive relationships between an aspect of one case and aspects of others. Such relationships are rarely brought out. Our hypertext systems capture these many lessons that are missed in strict case-by-case (or problem-by-problem) instruction. In an ill-structured domain, facilitating retrieval of multiple (partial) precedents is important, because understanding what to do in a given case context will usually require reference to more than any single prototype--the case in question will be "kind of like this earlier one, kind of like that one," and so on. Also, understanding of the case in question will require that various concepts be brought to bear and integrated; this, too, is facilitated by the multiple conceptual coding scheme employed in our systems.

There are several other benefits of the multiple-conceptual coding of multiple case segments. A power/efficiency advantage is that it allows the hypertexts to automatically generate large numbers of lessons (many "landscape criss-crossings"). If, for example, each of 20 cases is divided into an average of 10 case segments, each with a value on 15 relevant thematic dimensions, there is a many-fold increase in the number of possible automatizable instructional comparisons and contrasts that results from having 200 case segments (instead of 20 full cases) intertwined by relationships in the 15-slot vector.

Also, the use of case segments prevents the subsumption to a "common denominator" that occurs when larger structural units are used: An interesting local element of a case will tend to get lost if it has features that are not present in other parts of the case (when the monolithic case is the structural unit). Using small case segments (minicases) helps to retain the *plurality* of situations.

There is another virtue of the division into case segments and the multidimensional coding of the segments that relates to keeping case understanding from being overly simplified. In an ill-structured knowledge domain, by definition, there is sufficient variability across cases (due in part to the interaction of the many factors that make up complex cases) that the set of cases that might be nominally grouped together under some schema or classification will be greatly variable in their particulars. A case, instead of being represented as one kind of thing, conveying one kind of "lesson," is instead clearly shown to the learner to be *many things*. Cases of the same nominal type have different segments or scenes that are demonstrated not to be the

same, and each of the segments is shown to have multiple significances. Therefore, the common temptation to nest cases uniquely under a single superordinate conceptual category will be resisted, making it less likely that the complex relationships among cases in a domain will be artificially regularized. In an ill-structured domain, cases are related to many different concepts of the domain, and it promotes dysfunctional simplification to hierarchically nest or "slot" cases under single conceptual categories (e.g., "The following cases are examples of X [only]"). When there is considerable across-case variability, as there will be in an ill-structured domain, cognitive flexibility requires that case information be coded conceptually for the many different kinds of use that new situations may require.

The thematic coding scheme and the landscape criss-crossing system of instruction result in a weblike multiple interconnectedness on multiple dimensions that is not subject to the limitations of instruction that is characterized by a single organizational slant. Instead of a single text with a single organizational scheme and a single sequencing of comparisons and contrasts, our hypertexts allow the same information to be automatically reconfigured according to a huge number of possible organizational schemes, determined by using subsets of the multiple thematic coding space--our hypertexts enable the *virtually limitless automatic generation of new text configurations*. Because of the richness of ill-structured domains such as biomedical science, each of these text configurations teaches some case- (experience-) grounded lessons that would not have been taught (or easily seen if taught) from another text's organizational perspective. Such additional experiences and perspectives are always helpful in a complex domain--a physician never learns all that it would be helpful to learn (which is why additional experience is always valued in a physician). Hypertext systems designed like the *Cardioworld Explorer* have the potential to systematically *consolidate the process of acquiring experience*.

Yet another virtue of the multiple interconnectedness along multiple dimensions of the representations that our systems build has to do with the problem of reciprocal misconception compounding that we have observed in our studies of medical students and physicians (Coulson, et al., 1986; Feltoovich, et al., in press). Misconceptions bolster each other and combine to form seductively entrenched *networks of misconception*. Our approach helps to forestall the development of misconception networks by developing a kind of *positive reciprocation*. Because *correctly conceived representations with a high degree of multiple interconnectedness are established*, the fresh entry of fallacious knowledge at any node in the weblike network will fire off so many connections that it would be likely to activate some *misconception-disabling correct knowledge*. Before you can go too far wrong, you are likely to touch something that sets you right.

2.3.7. Active Participation, Tutorial Guidance, and Adjunct Support for the Management of Complexity. In an ill-structured domain, knowledge cannot just be handed to the learner. A priori codifications of knowledge are likely to misrepresent. (That is part of what ill-structuredness means.) Hence the importance, increasingly widely recognized today, of active learner involvement in knowledge acquisition, accompanied by opportunistic guidance by expert mentors (which can be incorporated in a computer program--it does not have to be live, one-to-one guidance). Furthermore, aids must be provided to help the learner manage the added complexity that comes with ill-structure. Our hypertext programs allow learners to explore complex conceptual landscapes in many directions, with expert guidance and various

kinds of cognitive support (e.g., integrated visual displays). When there are limits to the *explicit* transmission of knowledge, learners will need special kinds of help in figuring things out for themselves (see Barrows & Tamblyn, 1980; Collins, Brown, & Newman, in press; Spiro, et al., 1987).

2.4 Recapitulation: A Shift from Single to Multiple Representations and from Generic Schema Retrieval to Situation-Specific Knowledge Assembly

In general, we argue that the goals of advanced knowledge acquisition in complex and ill-structured domains can best be attained (and the problems we have identified avoided) by the development of mental representations that support *cognitive flexibility*. Central to the cultivation of cognitive flexibility are approaches to learning, instruction, and knowledge representation that: (a) allow an important role for *multiple representations*; (b) view learning as the multidirectional and multiperspectival "criss-crossing" of cases and concepts that make up complex domains' "landscapes" (with resulting interconnectedness along multiple dimensions); and (c) foster the ability to assemble diverse knowledge sources to adaptively fit the needs of a particular knowledge application situation (rather than the search for a precompiled schema that fits the situation). We suggest that theory-based computer hypertext systems can implement the goals and strategies of Cognitive Flexibility Theory, engendering multiple cognitive representations that capture the real-world complexities of the kinds of cases to which abstract conceptual knowledge must be applied.

Section 3.0 PATTERNS OF FAILURE IN ADVANCED KNOWLEDGE ACQUISITION IN COMPLEX DOMAINS

In this section of the proposal we will present the results of the first major part of the program of research, the studies of the nature of learning failure in advanced knowledge acquisition.

Chapter 3.1 The Deep Structure of Complex Domains and the Development of Misconceptions: The Nature of Conceptual Understanding in Biomedicine

3.1.1 Overview of Chapter and Introduction

A general framework is presented for studying the acquisition and cognitive representation of conceptual knowledge in a complex domain, biomedicine. This framework is also applied in analyzing the nature and development of misconceptions. Central to our approach is the selective and highly concentrated analysis of clusters of complex concepts, both as to their true nature and the manner in which they are understood by learners. What we find in these analyses is a widespread tendency for medical students to develop significant errors in conceptual understanding. These errors include specific misunderstandings of concepts and, more generally, maladaptive biases in the thought processes that are brought to bear in dealing with conceptual complexity. Particularly noteworthy is our observation of an insidious tendency for misconceptions to compound each other within a general climate of oversimplification, producing large-scale, mutually reinforcing, and durably entrenched areas of faulty understanding.

One such widely held fallacy concerns the causal basis of congestive heart failure at the level of the muscle cell. This fallacy and factors contributing to its acquisition and maintenance in belief are discussed in considerable detail in the paper. The implications of our findings for learning and instruction are discussed. The misconception is noteworthy because it serves as an intact example of the application of the theoretical framework which is developed in the first part of the chapter. The heart failure misconception brings into focus important commonalities we have observed in the nature and developmental pattern of a variety of biomedical misunderstandings (Spiro, Feltovich & Coulson, 1988; Spiro, Feltovich, Coulson & Anderson, in press; Feltovich, Coulson & Spiro, 1986):

1. *Multiplicity.* Many influences contribute to the acquisition and maintenance of misconceptions, some of which are associated with the learner, some with the educational process, and some even with the practices of biomedical science research.
 2. *Interdependency.* Overall misconceptions can be represented as *reciprocating networks* of faulty component ideas which mutually bolster each other and, in turn, support the overall misconception.
- ... *Oversimplification* of complex biomedical phenomena and concepts appears to be a major force in the acquisition and maintenance of misconceptions.

Thus, what may seem at first to be a simple misconception, easily describable in a single sentence, will turn out in fact to have numerous complexly interrelated layers of underlying meaning. It is in this sense that we refer to a *deep structure* of ideas. As will be made clear later, these features would likely go unnoticed without a concentrated analysis of the complexities of individual concepts and of the finer threads of relatedness that exist among neighboring concepts.

The chapter is organized into three main sections. In the first part of the chapter we highlight the central features of our approach and the advantages that accrue from its application. An in-depth analysis of the structure and genesis of the heart failure fallacy in the second part of the chapter illustrates more specifically the intricate and convoluted patterns we have found in the development of biomedical misconceptions. Finally, we discuss the implications of our findings for learning and instruction, and for the practice of medical education.

3.1.2 The Concentrated Study of Concepts and Conceptual Families in Advanced Knowledge Acquisition

Our approach eschews broad but superficial coverage of large numbers of biomedical concepts in favor of in-depth coverage of small sets of important and complex concepts. This decision is motivated by our interest in advanced knowledge acquisition, in how people come eventually (if ever) to understand and apply complex material well. The ideas of interest are ones that are intellectually challenging, ones that require much time and much effort to master; alternatively, they are ones for which, despite effort, understanding may never progress very far for various reasons (e.g., because a particular conceptualization acquired early in learning blocks progress). It is with regard to advanced knowledge acquisition for complex material that current theories of learning are most deficient and current educational methods least effective. In particular, although much has been described about experts and novices (e.g., Chi, Feltovich & Glaser, 1981; Larkin, McDermott, Simon & Simon, 1980), little is known about the *acquisition* of the advanced understandings found in expertise or about the best educational methods for fostering them (but see Collins, Brown & Newman, in press; Spiro, Vispoel, Schmitz, Samarapungavan & Boerger, 1987). Because the kinds of ideas we are interested in are difficult and complex and because the processes of learning and understanding them are also complex, and not well understood, the process of *studying* these matters must be detailed and comprehensive.

We believe that being able to deal with complexity is essential if those who are to use complex knowledge, for example, medical practitioners, are to be able to respond effectively to deviations from routine situations that require flexible and adaptive thought. The perspective of advanced knowledge acquisition and understanding, with its emphasis on the psychological management of complexity, is discussed in the first subsection below.

A second issue that arises from our approach concerns *concept selection*. Because we deal with a restricted set of concepts, we must ensure that those that we do address are important ones, central to the development of biomedical expertise, and ones that are in need of study because they present problems to learners. Our approach to concept selection, based on concept "nominations" by medical teachers and practitioners, is described in section 3.1.2.2.

In the section devoted to the kinds of generality derivable from concentrated conceptual analysis we discuss some of the important advantages that accrue from our approach. Besides the practical benefits that result from the study of concepts suggested by practitioners and teachers as being important and difficult, several kinds of general implications are derived as well. The results of our in-depth analyses possess several kinds of *generality*, including: themes related to content knowledge structure; themes of cognitive representation (in response to cognitive challenges of complexity); themes in the presuppositional, prefigurative schemes that are brought to bear on complex ideas; biases in the management of conceptual complexity; and patterns of misconception *development* that *recur* in different conceptual families, including two types of conceptual compounding. These forms of generality are observed both for incorrect knowledge about individual concepts and for higher-order misconceptions that evolve because of compoundings of several component misconceptions.

3.1.2.1 Complexity and Advanced Knowledge Acquisition

Certain aspects of our stance with regard to complexity are worth noting at the outset. These concern what we mean by complexity and what we argue is the appropriate place of complexity and the reaction to it in learning and instruction, especially at more advanced levels of knowledge acquisition in a subject area.

Aspects of conceptual complexity and domain ill-structuredness. Concepts can be difficult and complex for a variety of reasons. In general they are difficult because they make unusual demands on cognition as compared with more mundane cognitive practices or abilities. There are at least four categories of demand that make a concept hard:

1. *Unusual demands on working memory.* Included here are large numbers of nested steps or goals that must be managed, or large numbers of variables or simultaneous processes that must be reconciled--aspects of multidimensionality.
2. *Unusual demands on formal representation.* Two aspects are important here. One is the degree of abstraction necessary for understanding--the extent to which the concept includes components that are less concrete (e.g., ideas of *rate* or *acceleration* of flow). Another demand involves the semantic distance between concepts and their formal symbolic representations (e.g., as in equations). The number "2" and the concept of *two* that it represents are fairly close. In contrast, the concept of *stroke volume of the heart* and the "SV" that appears in equations involving stroke volume are quite distant. Coordination of the constraints on the symbolic formalization and the (often quite different) constraints on the real *referents* of the symbols can cause great difficulty.¹
3. *Unusual demands on intuition or prior knowledge.* Concepts can be at odds with ostensibly related prior experience (so that, for instance, importation of seemingly relevant analogies is detrimental), or they can be counterintuitive or discrepant with common sense.

4. *Unusual demands on notions of regularity.* Concepts can be ill-structured: They can be highly variable in their application, requiring tailoring to context, recognition of numerous exceptions, the ability to deal with substantial "grey areas," and so on. Concepts can be highly dependent for their understanding on other concepts they tend to co-occur with and interact with, requiring understanding of elements of reciprocation among families of related concepts.

As will be discussed, characteristics such as these cause problems partly because the demands they pose are at odds with more typical cognitive modes and instructional practices that may be adequate and appropriate for more well-structured or simpler ideas.

The perspective of advanced knowledge acquisition. Any content area will have facets of well-structured simplicity. These facets are dealt with quite well by existing theories of learning, instruction, and knowledge representation, and they will not concern us here (see Spiro, 1980; Spiro et al., 1987, for reviews). Other facets will be complex and ill-structured. Often, when complexity is encountered in instructional settings, it is transformed into a simpler form to convey an introductory overview of the area to beginners. Complexity is then introduced incrementally, along with directed challenges to earlier "scaffold" models, in order to guide learners toward greater sophistication (e.g., Glaser, 1984). At the early stages of learning, the motivation for strategies of simplification, that of providing nonconfusing immediate access to a subject area, can be appreciated.

However, we have two concerns about this approach. First, we have found that initially simplified approaches often impede the later acquisition of complexity; learners resist what is difficult when they have something simple and cognitively satisfying that has already been learned to fall back on, and early, simpler models provide "lenses" for filtering out ill-fitting aspects of the new material (see Felton, Coulson & Spiro, 1986; Spiro et al., in press). Second, we have also observed that the tendency to reduce complexity early is carried over as a kind of habit of thought or instruction to more advanced stages of knowledge acquisition; that is, the reductive modes of learning and instruction applied in initiation to a topic are also employed in attempting to attain a deeper, more sophisticated understanding.

Accordingly, we distinguish *advanced knowledge acquisition* from the kind of learning and instruction frequently advocated for beginners in a subject area (Spiro et al., in press). For example, it may be reasonable to consider the primary goals of instruction at early stages of learning to be that students avoid confusion and not be daunted. The perspective of advanced knowledge acquisition, on the other hand, assigns the highest priority to getting the concepts *right*, even if some difficulty is part of the cost (of course, difficulty should be reduced as much as possible, without sacrificing the *integrity* of the concepts).

Another factor that differentiates advanced knowledge acquisition from standard practices of beginner-level instruction is the nature of assessment: What it is that is considered appropriate and satisfactory evidence of learning. In the beginner modes, reproductive or imitative *memory* for key facts and definitions is the most frequently employed learning criterion. From the perspective of advanced knowledge acquisition, demonstrations of

learning involve the novel *use, application, or transfer* of explicitly taught material (subject, of course, to adequate memory for *prerequisite* information, for the rudiments of a subject area). *Memory* for material that has been taught is not the same as *learning from* instruction (Spiro, 1980). In fact, it seems likely that those factors that promote accurate memory (e.g., tightly compartmentalized, insular mental representation) are *antithetical* to the development of usable/applicable knowledge (Spiro, 1977, 1980; Spiro et al., 1987). In one recent study of students who had just completed a cardiovascular curricular block, we found no positive correlations between performance on memory-type tests and either tests of deep inferential understanding of course concepts or applications of those concepts in clinical tasks (Coulson, Feltoovich & Spiro, 1987).

It is important to emphasize that there is no clear split between early and advanced knowledge acquisition--the course of the latter is highly dependent on the former. Understanding is a process in which what is learned about a topic at one time will affect the kinds and levels of understanding that can be achieved on that topic at a later time. The impediments and, possibly, limits to the level of understanding that appear to derive from the strategy of artificially reducing inherent complexity in early learning raise concerns about this strategy. If the goal, ultimately, is deep understanding (ultimately "getting it right"), then the desirability in some situations of fundamentally different approaches to early learning is raised. Alternatives might include exposing students to the full complexity of ideas from the beginning, recognizing that their initial feelings of mastery, accomplishment, and satisfaction may suffer temporarily, but that their horizons of understanding may be greater. Another alternative, discussed at the end of this chapter, involves the instructional use of relatively simple and understandable pedagogical components, but with a guidance system that at every step of learning highlights the limitations and misleading aspects of these components, as well as the linkages and sources of mutual embellishment among them.

The distinctions we have drawn are important because one of our central claims about advanced knowledge acquisition for complex concepts is that the processes of learning, instruction, and mental representation entailed are often diametrically opposed to those commonly found in introductory learning. For example, in examining the heart failure misconception later in the chapter, the same biases found to be maladaptively reductive in the development of misconceptions would appear to be a model set of familiar prescriptions for teaching and learning (and even scientific theory building) in the context of typical introductory learning or in a simple and well-structured domain. What in one situation might be taken to be a goal, in the other is a main obstacle to goal attainment.

3.1.2.2 Methodology for the Empirical Selection of Concepts for Concentrated Study

An approach such as ours, which focuses in detail on students' learning and understanding of a small number of concepts, requires that the concepts chosen for study be well selected. In order to support our goals, the concepts should meet three criteria: (a) they must be *perceived* by the medical community as important to the practice of medicine; (b) they must be important to the practice of medicine; and (c) they should be *difficult* for students (who eventually become medical practitioners) to learn, understand, and apply.

Perceptions of importance matter, if only for pragmatic reasons having to do with potential receptiveness to our findings. It would be discouraging, after much work, to receive a response of "So what?" because a topic we have chosen is viewed as unimportant. *Actual* importance has theoretical significance. One of our goals is to be able to trace the effects of the learning and understanding of basic science ideas into clinical practice. These links can be subtle and circuitous, and our chances of being able to trace impact are greater for the important concepts. These concepts are ones that are critical as building blocks for understanding others, that apply in many ostensibly different circumstances, or that interlock greatly with other concepts. *Difficulty* matters for reasons of potential educational usefulness of our work; difficult ideas, of course, are the most problematic for teachers and students. Perhaps most importantly, it is the most difficult ideas that provide the best opportunity for studying advanced knowledge acquisition, how people adapt to and manage conceptual complexity.

For all of these reasons, we have adopted an *empirical* approach to choosing concepts that relies heavily on the judgments of medical teachers and practitioners. There is a trust that such practitioners have good intuitions regarding the concepts that are both difficult and important to the practice of medicine. We are studying concepts nominated by medical school faculty, both basic scientists and clinical faculty (many of whom also practice as physicians). This was done initially through polling of faculty at our own medical school and of a sample of community physicians, and has continued in a survey of all medical schools in the United States and Canada (Dawson-Saunders, Feltovich, Coulson & Steward, 1987). In the survey, faculty are asked to nominate concepts that meet the criteria mentioned earlier: Importance to the practice of medicine and, in view of teaching experience, chronic difficulty for students in understanding and application. Students' learning and understanding of concepts chosen in this way are then investigated in our laboratory.²

3.1.2.3 Kinds of Generality Derivable from Concentrated Conceptual Analysis

Obviously, there is a *prima facie* medical relevance to studying concepts judged by the medical community (in our survey) to be important and to have wide scope of application within medical practice. In addition, however, findings from the concentrated study of individual concepts have several other kinds of general relevance. We now sketch the various ways that what is learned from a close conceptual analysis has more general implications.

3.1.2.3.1 Generality due to the cross-contextual applicability of concepts: Content themes. The same basic knowledge and principles involved in one context may apply in other contexts. For example, the factors that contribute to opposition (impedance) to the flow of blood--resistance related to vessel diameter and blood viscosity, compliance having to do with vessel diameter and the stretchiness of vessels, and inertance having to do with vessel diameter and the mass or density of the flowing material--are also operative in pulmonary air flow. The latter two factors are particularly germane in flow systems in which the flowing material undergoes cycles of acceleration and deceleration (alternating current circuits) as happens in both the cardiovascular and pulmonary systems. Correct understanding in either system, of accelerative effects as they relate to oppositional factors, will yield beneficial transfer to the other system (and incorrect learning will produce negative transfer).

3.1.2.3.2 Generality in types of cognitive challenge: Representational themes.

Phenomena or concepts (which themselves represent phenomena) must be mentally represented. A mental representation captures some aspects of what is being represented but may omit others or construe aspects in various alternative ways. Cognitive representations can vary in their fidelity to what is represented, and one of the things that determines the use of one representation over another is the cognitive challenge or difficulty it imposes (see Siegler & Klahr, 1982).

We often find that the mental representations used by individuals show a consistent pattern of response to cognitive challenge. The same cognitive process demands that make one concept difficult are often found to make other, ostensibly different, concepts difficult. An example is the representation of phenomena involving numerous processes or variables. One of the things that makes understanding of opposition to blood flow (cardiovascular impedance) difficult is the need to reconcile simultaneously the effects of resistance, compliance, and inertance (the latter two through their contributions to reactance), compounded by their interactions with heart rate, blood density, blood viscosity, and various other factors. (Reconciling these same numerous variables is, of course, an important challenge in understanding pulmonary air flow.) In a completely different system, discerning the interactions among numerous variables and simultaneous processes is a major cognitive demand that makes the understanding of acid/base and electrolyte balance difficult. Anticipating a general issue that will be taken up in the section on "Themes of Complexity-Reduction," a response of students in both the cardiac and acid/base domains is to under-dimensionalize: The phenomena are represented in ways that constrain them to subsets of the operative variables, rather than dealing with the complex interactions. With regard to oppositional effects in blood flow, the numerous different factors are interpreted in terms of, or as being analogous to, resistance (Feltovich, Coulson & Spiro, 1986).

Another example, involving a different kind of cognitive challenge, is the representation of rates, as is necessary in understanding blood flow as a *rate* of flow (rate of change of position of volume). Such understanding requires the representation of a ratio of differences and poses substantial difficulty for students (even before consideration of rates, which is critical, for example, in understanding accelerative effects). A common student response is to represent *rate* of blood flow as blood *volume*. This representation is equivalent to treating blood flow as its mathematical integration. Understanding of rate is important in any circumstance in which any kind of material flows, so that a finding of rate-related misrepresentation in one instance has general ramifications.

3.1.2.3.3 Generality in the importation of explanatory models: Themes of field prefiguration.

Prior to selecting the specific cognitive representation that will be applied to a phenomenon or concept, and constraining the nature of that representation, a much more fundamental and often unnoticed issue has often already been decided: What general form should an explanation take? Questions of this sort draw heavily on individuals' or disciplines' beliefs about the fundamental nature of the phenomena they deal with. A prefigurative scheme is a kind of "world view," a set of (often tacit) presuppositions about the fundamental nature of the world (or some circumscribed domain) and about what constitutes legitimate evidence and explanation. Prefigurative schemes are subtler, although perhaps more powerful in their influence on thought, than conceptual content or even its particular cognitive representation (discussed

above). Both of the latter, in a sense, refer to things one thinks *about*; in contrast a prefigurative scheme is something one thinks *in terms of* (cf. Bransford, Nitsch & Franks, 1977), a kind of lens that one sees with and that at the same time determines what is excluded from view. Such schemes resemble what have been characterized within science as "paradigms" (Kuhn, 1971). Philosophers of science may study and be aware of alternative paradigms held by the scientific practitioners of a time, but the practitioners themselves function *within* them (*tacit* knowledge); similarly for the implicit paradigms that guide the thought of a medical student.

The notion of a prefigurative scheme is perhaps most easily illustrated in the area of history. There are a variety of established explanatory models that are utilized in historical accounts (cf. White, 1973). Some historical explanations will presuppose a historical *mechanism*, with one event causing another, within chains of cause and effect sequences of events. Other historians will import the model of an *organism*, seeking out some teleological principle of historical development that unites many diverse historical events, but not in a cause-and-effect manner; the historical meaning of an event is treated as more than the sum of its parts or their effects. Still other historians will provide historical accounts that resemble *networks*, with an explanation accumulating as a function of the compounding of multiple interrelationships across the sectors of historical space. Once the metaphor of a mechanism, or an organism, or a network has been adopted as an underlying model, the nature of an explanation, the kinds of phenomena to be selectively focused on, and other important issues have to a great degree already been determined (Pepper, 1942).

Such broad-scale points of view can be applied to biomedical phenomena as well. For example, under a mechanistic scheme, in understanding or explaining biological processes one is more likely to look for causal agents and acts, and to decompose processes into pieces, steps, and causal chains. Under a more organicistic view, the process as a whole will be taken to have primacy over its parts; the function of the process in supporting some larger objective of the organism (e.g., homeostasis) provides guiding principles for the functioning of components, and analytic decomposition is accomplished only at some loss.

When prefigurative schemes are engaged, they have at least three kinds of effects. First, they affect thinking, for example, by affecting what from a situation will be represented cognitively and what will be excluded (or just not noticed). Second, they provide general rules for what constitutes legitimate explanation of, or evidence for, what is going on in a situation. Third, they reflect fundamental beliefs about the *nature* of the world and its phenomena.

Prefigurative schemes can be applied at many levels. At the most encompassing level are ontological schemes of the sort we have been discussing. These address such basic issues as the relationship of parts to wholes. On a still very broad scale, but more concrete, are schemes like the geocentric theory of the universe. Broad-scope scientific theories, such as Newtonian mechanics, also function as lenses with their own explanatory schemes, modes of evidence, and so on. Schemes for the interpretation of more localized phenomena can likewise come to capture a field and its practitioners in such a pervasive way that they become presumptive, tacit, and prescriptive. Almost as second nature, they come to *be* the phenomena they were originally designed to characterize. With regard to the heart failure misconception to

be discussed, the Sliding Filament Theory of muscle contraction will be seen to play a role of this sort.

The discussion thus far of prefigurative schemes should not be taken as implying that only one such scheme can be held or applied by a person at the same time; and, we do not wish to suggest that there is some best lens. For one phenomenon, one scheme will be better, for another phenomenon, another scheme will be better, and multiple schemes applied to the same phenomenon may compensate for the way each single scheme selectively accentuates particular aspects. Rather, we simply wish to suggest the existence of schemes that prefigure cognition and to discuss something of their nature and consequences for cognition. Misconceptions can result from the engagement of a prefigurative scheme that is faulty, inappropriate, or inadequate (e.g., one whose selection of aspects of a phenomenon to be understood is ill-fit to that phenomenon). Also, the engagement of prefigurative schemes can produce errors in systematic ways (e.g., misunderstanding of continuous aspects of processes by the importation of a mechanistic scheme that segments those processes).

These commonalities in the ways that a prefigurative scheme can induce misconceptions (i.e., prefigurative themes) are yet another aspect of generality that can accrue from the in-depth study of concepts. Concentrated conceptual analysis can reveal the *underlying* model that is guiding the more specific aspects of a concept's representation. By close analysis of several concepts, patterns of misconception development induced by prefigurative schemes can be identified.

A study of individual differences in prefigurative belief systems lent support to the hypothesis that learners differ in fundamental ways, along the lines of adherence or eschewal of an analytic, integrative, reductive world view. Sixty medical students were asked a large set of questions related to their underlying assumptions about the nature of phenomena in the world and about the nature of learning and understanding. Each question had a reductive and a nonreductive alternative. Although the questions covered highly diverse arenas, a factor analysis yielded as its primary factor a cluster of 15 disparate items, all perfectly aligned along the hypothesized dimension of Simple versus Complex. That is, some individuals have very consistent, across-the-board reductive world views while others prefigure the world and their understandings just as consistently with a view that acknowledges complexity, irregularity, and so on.

3.1.2.3.4 Generality of cognitive biases in the management of difficult concepts: Themes of complexity reduction. So far we have seen three kinds of generalization associated with the close analysis of sets of concepts: Content themes, representational themes, and prefigurative themes. In this section we show that all of these themes are associated with a directional bias in the way they are cognitively instantiated. The bias in each case involves the reduction of complexity.

"It costs energy to make surfaces that are rough" (Gleick, 1987). A sphere is the lowest energy configuration that a soap or water bubble can assume, and it is the one toward which its surface forces tend; to make a bubble assume any other configuration requires the input of energy. The preceding quote is from a discussion of the formation of a snowflake as a process that at all stages must overcome the surface tension of water, which would tend

naturally to make the snowflake round. Everywhere in nature there are strong forces tending toward simplicity. Recent theories of information processing and knowledge representation (Smolensky, 1986) have suggested a similar tendency for cognitive processes and interpretations to settle into the lowest energy configurations, in a sense the ones that can be assumed with the least effort. Hence, it should come as no real surprise that there are such simplifying forces playing a role in the understanding of complex concepts; as in the making of an exquisite snowflake, intricate understanding may also have to oppose such "reductive forces."

Thus, another area of generality that emerges from the study of concepts involves the nature of and principles governing these kinds of cognitive reductions, their ramifications for misconception, and, ultimately, insight into ways they can be overcome in advanced knowledge acquisition. The concepts we have chosen to study are difficult and complex. Close conceptual analyses reveal the strategies (intentional and unintentional) that students and teachers adopt in the face of that complexity. Students are sometimes able to cope with difficult concepts without trivializing them, but, unfortunately, this does not appear to happen very often. Instead, there is a proclivity toward the strategic mismanagement of complexity, involving various forms of *oversimplification*. Close analysis of conceptual understanding in a variety of biomedical domains has revealed numerous misconceptions that have in common some form of convenient and spuriously supportable oversimplification of a complex concept or phenomenon. We have referred to these comfortable oversimplifications as *seductive reductions* (Spiro, Feltovich & Coulson, 1988). Just considering the use of analogy in learning, we have observed biomedical misconceptions resulting from eight different varieties of spurious reduction of a new concept to an analogical source (Spiro et al., in press).

The misconceptions that we have characterized as seductive reductions result from the application of a large set of beliefs, motives, and cognitive operations that we refer to individually as reductive biases and collectively as *The Reductive Bias* (Coulson, Feltovich & Spiro, 1986; Spiro, Feltovich & Coulson, 1988). Many of these reductive biases will be illustrated in the context of the heart failure misconception. For now, we will give a few illustrative examples, concerning, first, reductive biases in conceptual understanding and, second, such biases in learning. It will be noticed that both of these kinds of reductive biases can operate in connection with all three kinds of themes we have discussed: Contentive, representational, and prefigurative (see the subsection below, "Realms of Reduction").

Conceptual biases involve systematic ways of spuriously reducing complexity in understanding. (Note that despite some similarities, conceptual biases are not the same thing as biases in judgment and decision making; see Elstein in press, Kahneman, Slovic, & Tversky, 1982, for discussion of the latter literature.) For example, there is a tendency in understanding the cardiovascular system to eschew dynamic (continuous, changing, etc.) interpretations in favor of a more static view, as when students equate blood flow with blood volume (*static bias*--a conceptual bias applied to a representational theme). Other examples of conceptual biases include the following:

- *Step-wise bias*. In another conceptual bias involving representation, continuous processes are broken down into discrete steps, with loss of properties that exist at the holistic level. This is often done by students in representing the continuous flow

of blood as a sequence of steps, causing them, for example, to misunderstand *relationships* between outflow and inflow to the heart.

- *External agent bias.* Intrinsic characteristics of entities or processes are attributed instead to external influences--as when vascular compliance is taken by students to be the vessel's ability to respond to orders from the nervous system. This confuses external mechanistic forces with inherent, internal processes of development (organic processes), a common consequence of applying mechanistic *prefigurative* schemes.

- *Prior analogy bias.* New concepts are interpreted through already held (often simpler) models, often imported from extra-instructional experience (e.g., the cardiovascular system is interpreted as being too much like household plumbing). This is a bias toward assuming more *contentive* similarity than is warranted.

- *Common connotation bias.* Technical terms are interpreted according to their everyday, common language meaning (e.g., the different types of erroneous models of vascular compliance we see correspond to various dictionary definitions of "compliance," as in the example above--also see Feltovich, Coulson & Spiro, 1986).

- *Restriction of scope bias.* General physiological principles are thought only to apply in specific instances (e.g., within conditions of pathology). This bias, in caricature, might support the proposition that the laws of gravity hold only on earth. In the biomedical arena, vascular resistance is taken by some students to be a property only of small blood vessels.

Acquisition biases are modes of addressing complex ideas during learning that attempt to make the ideas more tractable. Examples include *under-dimensioning*--a *representational* approach of teaching or learning multivariate phenomena one dimension at a time, with a goal toward eventual dimensional reassembly; *atomization and extirpation*--a *prefigurative* bias that extracts and isolates components from a multi-component system, with the assumption that their behavior in isolation will faithfully reflect their behavior in context; and *sanitizing*--a *contentive* bias that involves focusing (in learning and instruction) on the clearest, cleanest instances of a concept, those most insulated from contextual effects, with the idea these will be representative of, or will be a "bridge" to, the messy contextual exceptions.

Note that some acquisition biases directly parallel and reinforce conceptual biases. An example is the approach to learning about continuous processes that involves breaking them into component processes with sequential steps and effects (as is done routinely, for instance, in the learning/teaching of the cardiac cycle--paralleling the step-wise bias).

Realms of reduction. The Reductive Bias applies to all three kinds of themes that we have discussed--those involving content, representation and prefiguration--and in each of these areas functions as a kind of "selector" that opts for simpler over more complex modes of interpretation.

With regard to *content*, the most common result of the application of The Reductive Bias involves treating things, or aspects of things, as being more similar or more stable in their characteristics than they are. (A consequence of this similarity bias is a tendency to produce overgeneralizations.) Forms of this contentive reductive bias appear in the use of analogy (Spiro et al., in press), e.g., *analogy treated as isomorphism*: analogy is treated as being the same relations, based on the same underlying mechanisms. Other examples include the following: *extension of attributes*--if A is like B with regard to attribute X, then A will be like B with regard to attributes Y, Z, and so on; *homogeneity of components*--explanations that account for one component of a system will account for others of the same nominal type, for example, "a muscle is a muscle," whether it pumps blood (cardiac) or "pumps iron" (skeletal); and *reduction of technical meaning to common connotation*, discussed earlier.

Many of the options available for cognitive *representation* can, likewise, be seen to span a simpler-to-more-complex dimension. Should a concept or phenomenon be represented cognitively as: unidimensional vs. multidimensional; static vs. dynamic; compartmentalized vs. interconnected; linear vs. non-linear; continuous vs. discrete; and so on? The operation of The Reductive Bias selects for the simpler pole among representational options, even when it is not appropriate.

There are ways that *prefigurative schemes* (lenses) can be characterized so that they too can be seen to vary in complexity. Such a characterization has been proposed by Pepper (1942) and involves two bipolar dimensions of classification: analytic-synthetic and dispersive-integrative. For analytic schemes, parts are the facts of a phenomenon, and combinatorial interaction of parts and synthesis are derivative; in contrast, for synthetic schemes wholes and emergent properties are basic and decomposition is derivative and reductive of the phenomenon. While this first dimension deals with the relative primacy of parts and wholes, the second addresses inherent orderliness. Dispersive schemes accentuate the irregularity and ill-structuredness of events, while integrative schemes presuppose more coherence, orderliness, classifiability, and so forth.

The Reductive Bias leads to "pole migration" with regard to these schemes, that is, to the adoption of those prefigurative schemes that are simpler and easier to manage cognitively--the analytic and the integrative. This manifests itself in reductive biases such as *atomization and extirpation*, the presupposition that "parts" extracted from context will faithfully reflect their properties and dynamics of operation in full context (analytic), and *uniformity of explanation*, the presupposition that complex processes are *really* governed by some single principle or mechanism (integrative).

The adoption of the *simpler* schemes promotes a kind of double jeopardy. First, it bolsters the application of *single* lenses or schemes to phenomena--a uniformity of 'lens' application. For example, dispersive schemes are antithetical to the notion that any uniform explanation can be applied everywhere, while this uniformity is the *forte* of integrative schemes. In turn, when analytic or integrative schemes are adopted, the world is then selectively interpreted that way, filtering aspects that fall "off-line," and bolstering the adoption of the reductive schemes in the first place. An integrative scheme promotes single lens application--single lens application bolsters the apparent efficacy of the integrative scheme. A

danger is the belief that the scheme is accounting for more of the variance in a phenomenon than it actually is. An example that plays a big part in the heart failure misconception is the assumption that the Sliding Filament Theory accounts for all muscle cell contractile dynamics, when, in fact, it accounts for only some.

Why are reductive cognitive modes of all kinds so readily adopted? One reason has to do with cognitive ease. It is easier to think that all instances of the same nominal concept, for example, *compliance*, are the same or bear considerable similarity (*contentive* reduction). It is easier to represent continuities in terms of components and steps (*representational* reduction). It is easier to deal with a single principle from which an entire complex phenomenon "spins out" than to deal with numerous, more localized principles and their interactions (*integrative prefigurative* reduction).

Furthermore, reductive tendencies are reinforced by many educational practices, particularly those associated with introductory teaching and learning. We have argued that the educational process in general is riddled by an implicit and unintended *conspiracy of convenience* (Spiro et al., 1987) to treat complex subject matter as simple. This makes it easier for students to learn, for teachers to teach, for textbook authors to write, and for testers to construct and grade their tests. Thus, besides making a difficult situation easier for everybody, The Reductive Bias finds the grounds for its acceptance already prepared, because the elements of this conspiracy have already been established in what are commonly taken to be effective modes of teaching and learning: highlights, conceptual "coat racks," sanitizing (not confusing students with all those exceptions), glossed exposure, etc.

Adding to the problem is that for some circumstances and some kinds of learning, such reductive maneuvers may be appropriate and effective. Much that has been pointed to negatively here is often taken as efficient and effective cognitive processing (e.g., *extension of attributes* as "default assignment"--see Footnote 2), effective representation (e.g., *decomposition* of a complex problem into components), or even standard canons of science (e.g., investigating the dynamics of complex phenomena by varying one dimension at a time). Such common practices are appropriate where they are appropriate--generally with regard to well-formed concepts and well-structured domains. However, in more complex and ill-structured domains these very same practices are problems, not principles. One might argue that what makes truly difficult concepts hard is that they fall "off-line" in ways that undermine mundane cognitive and investigative tools; if they were routine and usual, they would not be hard.

Generality of patterns of development of higher-order misconceptions: Themes of idea-compounding and spreading misconception. Although we focus at first on individual concepts, the fact that many of the concepts we choose form groups of conceptual clusters or families allows patterns of higher-order conceptual representation to emerge. We have discussed how our approach to the close analysis of individual concepts informs both about the detailed nature of the understanding of those important concepts and, by allowing observation of themes that recur across individual concepts, about certain general principles and problems of conceptual understanding. We now want to demonstrate a different kind of benefit: By the concentrated study of the individual concepts in a conceptual family, otherwise hidden patterns become evident in which *misconceptions compound each other* and thereby contribute to higher-order misconceptions. These higher-order misconceptions may be even more seductively

entrenched than the local ones. *Spreading activation* has been a popular topic in recent cognitive science (Anderson, 1983; Collins & Loftus, 1975) by describing the processes by which ideas compound each other, we introduce the notion of *spreading misconception*.

Such higher-order misconceptions take a variety of forms, of which we will illustrate two. One involves *compounding*: Misinterpretations of fundamental ideas can cohere in systematic ways such that belief in one makes belief in others easier. Consider the following example (see Feltovich, Coulson & Spiro, 1986 for a detailed treatment). The effect of vascular compliance (a factor related to the stretchiness of vessels) in blood circulation is often interpreted by students as a kind of "surrender," as enabling vessels to "give way" and accommodate increases in blood flow. This is in contrast to the *active* role of compliance (and elasticity, its inverse) in promoting the movement of blood. Believing compliance is primarily accommodative makes it easier to believe that veins, highly compliant vessels, are reservoirs, places for holding blood, as opposed to active participants in ongoing dynamic flow. And, *reciprocally*, the latter belief bolsters the former one. Furthermore, pressure, in relation to blood flow, is interpreted as hydrostatic "bucket" pressure and not in terms of dynamic gradients--consistent with the first two beliefs. Downplaying the active role of compliance bolsters belief that impedance (opposition to blood flow) is the same as resistance, ignoring or subsuming effects of compliant and inertial reactance (which, unlike resistance, are totally dependent on dynamic aspects of blood flow, as they interact with compliant properties of blood vessels and other factors). This restriction of oppositional effects to resistance (as well as most of the interpretations listed above) is consistent with a view of the cardiovascular system as a "direct current" as opposed to "alternating current" circuit--which is bolstered by the non-recognition of accelerative properties of blood circulation due to pulsatile (continuously changing) pressure. It is also consistent with a view of the vasculature (e.g., veins) as simply accommodating active effects from the heart. And so on.

The compounding and mutual bolstering of these individual interpretive units leads to a second kind of higher-order conceptual effect--*pervasive coloration*. All conspire to promote a *passive* coloration in the "view" of the cardiovascular system, especially regarding the vasculature itself. Once such colorative effects emerge, they can then feed back and shade further learning and interpretation.

The compounding of misconception is aided, in part, by the reduction of phenomena to the simplified poles, for example, the representation of pressure statically and the contentive reduction of all impedance factors to resistance. Convergent shearing away of complicating factors provides the ground for individual factors, thus stripped, to complement each other and to compound. In turn, the broad colorations that then pervade the overall image of a phenomenon (e.g., the "passive" cardiovascular system) can turn back and reinforce thinking about individual concepts that way.

Generality of learning scenarios that cause conceptual error: The "anatomy" of misconception and its development. We have discussed the implications of close conceptual analysis under four thematic headings: content themes, representational themes, themes of field prefiguration, and themes of complexity-reduction. Each of these four general kinds of cognitive phenomena is associated with its own distinctive form of contribution to misconception and participates in spreading or compounding the development of higher-order misconceptions

(a fifth kind of theme). Considering together the first five kinds of generality described in this section, the result of our approach to close conceptual analysis is a detailed *anatomy* of the nature of conceptual understanding and the development of conceptual error, at the levels of the individual concept and conceptual clusters. This includes insight into the nature of misconceptions, their sources and their interrelationships.

The investigation of concepts in detail, but concepts that are themselves interrelated (a cluster approach), can uncover such interdependencies. This cluster approach enables investigation of emergent effects and broader scale misconceptions that would not be possible through the investigation of the constituent concepts in isolation. As the example of the misunderstanding of congestive heart failure will illustrate, the sources of misunderstanding can be many and can run far below the surface misunderstanding itself. The specific nature and causes of the heart failure misunderstanding to be discussed would not have emerged without our detailed and interrelated probing of such important constituent concepts of cardiovascular function as: *opposition to blood flow* (impedance), which pertains to factors which oppose the flow of blood; *the Frank-Starling cardiac function* and *Guyton's vascular function relationships*, which capture the intrinsic control of cardiovascular flow, through the interlocked regulation of vascular function by the heart and of cardiac function by the vasculature; *cardiac muscle activation and control of contraction*, which pertains to how the heart muscle itself works to produce its pumping action; and to a lesser extent, *cardiac hypertrophy*, *cardiac electrophysiology*, and *energetic metabolism*.

Only by looking at these interrelated concepts together, following in-depth investigation of each, could the pattern of interlocking misconception that we will present be detected. In other words, a broad and superficial probing of heart failure knowledge would have missed much of the important underlying basis for the misconception. In a related way, broad, superficial or compartmentalized instruction on heart failure, by missing detailed treatments of constituent basic science concepts and their complex interactions, would result in the "climate of oversimplification" that produces misconceptions of the kind we have found.

In the remainder of the chapter, the misconception of heart failure is addressed as a more concrete example of the application of frameworks that have been introduced in these early sections. It will be shown how misconceptions interact to bolster each other, compound to yield yet higher-order misconceptions, and align in gross aggregation to produce pervasive colorations--all within an erroneous belief-structure about the basis of congestive heart failure. The various kinds of compounding of interactions suggest a network structure for the overall belief system, a network that yields synergistic strength greater than the component parts. Paralleling the learner's interpretations and modes of thought (related to addressing complexity) that lead to and support misconception, similar reinforcing influences from the instructional process and from some practices of biomedical science are also discussed. It is argued that rather than being external to the cognitive network of misconception and simply contributory, these "external" influences are entwined inseparably within the network itself; they are partly responsible for the component misconceptions that come to make up the network and are partly responsible for the patterns of thought and practice that legitimize and, hence, help to maintain the network. Learning, educational practice, and laboratory science are seen to be subject to the same reductive biases and, therefore, to be symbiotically enmeshed.

3.1.3 The "Anatomy" of a Misconception of Heart Failure

The misconception of heart failure is taken up next, following the introduction of some necessary background information regarding congestive heart failure and cardiac muscle cell function.

3.1.3.1 Background Regarding Heart Failure

Congestive heart failure is a syndrome or constellation of effects in which the heart loses its effectiveness as a pump and fails to maintain flow rates consistent with the needs of the body. The misconception to be discussed has to do with the dynamics at the subcellular level of the heart muscle that account for heart failure. In order to understand the nature of the misconception, a brief overview of these subcellular dynamics is necessary.

The heart pumps by contracting its myocardial muscle. In heart failure this contraction is inadequate, lacking sufficient force, shortening capacity, or both. Figure 1 depicts the subcellular part of a muscle, the sarcomere, which ultimately produces contraction.

Within each sarcomere are myosin and actin filaments. The sarcomere is the unit between the two vertical bands (Z-bands) depicted in Figure 1 (also depicted in the figure are adjacent sarcomeres; i.e., sarcomeres are aligned in series, abutting at shared Z-bands). The myosin (see Figure 1) is a fixed-length filament that has as components a number of little "arms" (cross-bridges) that ultimately produce contraction. Actin are other fixed-length filaments that are attached to both ends of the sarcomere. These structures, attached to opposite ends of the sarcomere (to the Z-bands), are pulled toward the center of the sarcomere (see Figure 2) during a contraction, shortening the whole sarcomere without substantially changing the length of either (myosin or actin) filament structure. These subcellular structures and dynamics, actin and myosin filaments of fixed length which "slide over" each other during contraction without change in length to either, are components of the Sliding Filament Theory of muscle contraction (Gordon, Huxley & Julian, 1966).

The cross-bridges of the myosin produce the force of contraction. During a contraction, the heads of the cross-bridges bind to available sites on the actin, pull one stroke toward the middle, release, rebind to the next available outward site, pull again and so on, progressively pulling the actin (which is attached to the outward ends of the sarcomere) toward the middle, resulting in shortening of the sarcomere. This process (binding, pulling, releasing, rebinding, pulling, etc.) continues, in the presence of adequate metabolic (energetic) materials available and necessary to drive the process (energy is needed to make the sarcomere "go"), as long as the contractile process remains active.

The *force* of contraction is a product of two main things. The first is *anatomical/mechanical* and involves the degree of alignment of actin surfaces with myosin "arms." One can see that if, somehow, the actin were to get stretched too far outward, some of




Figure 1. Schematic representation of the sarcomere.




Figure 2. Schematic representation of sarcomere during shortening.

the arms would have nothing to bind to and force would decline. This is partly due to the constancy of length of both the actin and myosin filaments during contraction, consistent with the Sliding Filament Theory. The second factor affecting contraction force is *activational*. It involves the number of arms that are recruited to participate (to pull) in the contraction (not all need be), and is regulated by the degree of energetic activation of the actin filaments' cross-bridge binding sites. One can envision the actual *number of arms* engaged in the pulling operation being activationally regulated, resulting in degrees of strength of the pull.

3.1.3.2 Preview of the Misconception

The misconception held by medical students, some medical practitioners, and portrayed in some medical texts holds that heart failure happens because the heart gets too big, stretching component sarcomeres so that actin does not overlap the myosin optimally. This results in fewer potential binding sites, with a corresponding reduction of force of contraction. Hence, of the two main factors affecting force production in a muscle, one anatomical and the other activational, in the misconception the anatomic, mechanical component is invoked, with general neglect of the activational component. In reality, it is the activational mechanisms that are germane to heart failure and the mechanical/anatomic component can play no role. This is discussed after consideration of a few representative statements reflecting the basic misconception given by subjects in our laboratory and in a textbook of cardiology. The quotes from subjects are from responses to open-ended but focused discussion questions (organized by the 'target' concepts, e.g., *muscle activation and control of contraction*, listed in the last section) about the function of the heart and cardiovascular system (see Footnote 1). The first quote is from a second-year medical student (Figure 3); the second, from an established cardiovascular physician (Figure 4); and the third, from a respected and reasonably current textbook of clinical cardiology (Figure 5).

All of these descriptions portray a common mechanism as explanation for heart failure: If the heart gets too large, then the individual sarcomeres contained within the myocardial muscle of the heart likewise get stretched; potential binding sites on actin structures are lost due to non-alignment; and hence, the force generated within a pumping stroke of the heart is diminished, producing heart failure. This is a mechanical, anatomical account of heart failure. While this account has a certain seductiveness--it *seems* plausible, and it invokes the usually well-learned Sliding Filament Theory--it is inconsistent with the best available physiological and pathological evidence for the basis of heart failure.⁵

In contrast to this account, the basis of heart failure is activational. The muscle in a failing heart is "sick" muscle, and the basis of this sickness is metabolic and energetic (Katz, 1977); the failing heart does not get energized sufficiently to do its job. While some details of this activational account remain to be specified within laboratory science, the mechanism of heart failure likely involves some combination of actual *damage* to sarcomeres, defects of biochemistry in the activation system itself, altered activity of ion pumps, and perhaps many other defects of muscle chemistry yet to be discovered (Nair, Cutilletta, Zak et al., 1976; also, see Coulson, Feltovich & Spiro, 1986, for additional related references).

Okay, discuss factors which cause the muscle contraction to be inadequate (in heart failure). Okay, this takes me back to the Frank-Starling mechanism (Note: The Frank-Starling relationship describes how the heart produces greater force of contraction as it is filled with greater volumes of blood. This relationship and its role in the heart failure misconception are discussed in greater detail later.) where there's a volume overload on the heart and this, um, leads to, um, the muscle cell spindles to be spread apart so there's very little overlap between the actin and the myosin and ah, the way I conceptualize this is with the cross fibers as sort of like little guys rowing or else people pulling on a tug of rope, and ah, with less overlap they're able to ah, develop less tension and the entire muscle, muscle fiber itself, the heart is able to, is not able to do its job as effectively...I like to think of the Frank-Starling as falling over the edge of where, with, as the further the muscle cell gets stretched, the less it's able to do its job effectively...I use the sort of the, conceptualize the number of rowers or people pulling on a tug of rope to explain that with increasing stretching of each cell they are less able to pull and increasing stretching of each cell they are less able to pull and shorten. And also, in the Frank-Starling mechanism where it's gone over the edge of where more lengthening of the cell doesn't help with tension.

Figure 3. Quote from a second-year medical student.

What the Frank-Starling relationship indicated was that ah, this ah, matter of ah, a muscle fiber being stretched, there comes a point at which ah, the stretching is no longer productive and, therefore, we tend to talk about the ah, healthy part of the curve, that's one phrase which is used commonly, which is the physiological ah, range that ah, as one stretches the muscle fiber, it will contract more forcefully. However, a point is reached ah, which is so-called plateau point where stretching the muscle fiber a little bit more is tolerated but does not result in any increased ah, contraction and force of contraction. And then if one stretches (the muscle fiber) a little bit further than that, we then go on to what is referred to as the down-slope of the curve where actually the ability to contract is decreased. The concept we have is that there is a point at which the left ventricle can be so dilated that it is no longer functioning.

Figure 4. Quote from an established cardiovascular physician.

Dilatation of the heart is also a factor that can decrease cardiac efficiency. The force of contraction of a muscle (including cardiac muscle) depends on the initial length of the muscle sarcomeres. When the sarcomere is initially stretched, this is associated with a more forceful contraction. The optimum sarcomere length is $2.2\mu\text{m}$. At this length the overlapping between the actin and myosin filaments is ideally situated to allow the cross-bridges between them to pull the actin filaments inward during contraction. When the sarcomeres stretch beyond this point, the actin and myosin filaments do not overlap so much and the cross-bridges cannot pull the actin filaments inward adequately. As a result, the force of muscle contraction decreases. This is the structural basis for Starling's law of muscle contraction.

Figure 5. Quote from a reasonably current textbook of clinical cardiology (Goldberger, 1982)

3.1.3.3 A Pervasive and Powerful Misconception

Despite its inappropriateness, the "mechanical overstretching" account appears to be widely held in the medical community. This account has been given in our laboratory by medical students and by some established cardiovascular physicians, and is often described in textbooks. As medical school faculty involved in the normal teaching activities of the school, we have seen this explanation commonly conveyed to students. In our laboratory work with medical students, explanations of heart failure as resulting, exclusively or in part (interspersed with other mechanisms), from mechanical overstretching of muscle fiber were given by 18 of 28 (64%) first- and second-year medical students (first-year: 7 of 14, 50%; second-year: 11 of 14, 79%).^a

In addition to its pervasiveness, the misconception appears to have a certain kind of insidious power. In much of our other research with medical students, we have often been able to trace a direct influence for misconceptions and errors to the primary instructional materials used by the students in their course work (Feltovich, Johnson, Moller & Swanson, 1984). Medical students take their assigned materials seriously; to some extent, "they are what they read." However, there is something about the heart failure misconception that undermines this general rule-of-thumb. The main assigned cardiovascular textbook (Katz, 1977) used by our medical student subjects contains a contemporary and appropriate activational account of heart failure; in the chapter on heart failure there is no mention of overstretching and sarcomere disinterdigitation as a basis for heart failure (although, as we have seen, students may read such accounts in other places). This overriding of primary instructional material emphasizes both the aggregate power of the misconception and the possibly circuitous nature of its influences.

As we have described, heart failure is activational and biochemical. Yet, students believe its basis to be mechanical and anatomical. Our question is, *Why?* We have already suggested that the reasons are many, and, at least in part, circuitous and subtle. In the next section, the question of why students (and others) acquire and maintain the erroneous belief, and the nature of the belief itself are addressed in greater detail.

3.1.3.4 The Nature, Acquisition, and Maintenance of the Misconception

In this section we investigate why the *mechanical overstretching* misconception of heart failure is acquired by students, the influences that contribute to the robustness of the belief, and the structural form of the misconception itself. As a preview to this section, four general points can be raised:

1. In the misconception of heart failure, dynamics of individual, isolated *skeletal* muscle fibers are mapped to the intact functioning of the whole heart (which, of course, is not skeletal and is a system of *many* fibers).
2. Properties of the whole heart that exist because it *is* a system containing complexes of fibers are neglected within the misconception.

3. The overall misconception contains *component* misconceptions that interact and support each other in reciprocating ways, forming a network. The interlocking nature of these components yields an overall structure, the misconception, that is itself highly robust.

4. Simplification of complex ideas, on the part of the *learner*, the *educational process*, and, in some instances, the *practices of laboratory biomedical science* contribute to the development of the misconception and support its maintenance in belief.

This section is organized around four component misconceptions that are involved in the overall misunderstanding. As overview, these four components are: (a) that the means by which an individual, isolated (skeletal) muscle fiber produces different contracting force at different lengths of stretch (the *length-tension* relationship of muscle fiber, L-T) is solely mechanically and anatomically based; (b) that an individual cardiac muscle fiber is like an individual skeletal fiber; (c) that the increased force that an intact heart (*in vivo*) develops when it is filled with more blood volume (and gets bigger) results from the same mechanisms that enable an individual fiber to produce more force (up to a point) as it is stretched to greater length; and finally, (d) that when a heart becomes too big it loses its ability to generate adequate force, hence fails, for the same reasons that an individual skeletal cell loses its ability to generate tension when it is stretched too far (that is, because of disinterdigitation of actin/myosin filaments). These component ideas are taken up, in turn, along with the introduction of some necessary background for each where necessary. For each component, educational influences toward its development on the part of the learner and the teaching process are discussed, as well as pertinent aspects of laboratory biomedical science.

Component 1: The Entire Length-Tension (L-T) Relationship for an Isolated Skeletal Muscle has an Anatomical Basis (*Length-Tension is Anatomical*)

Background to component 1. The length-tension relationship of a skeletal muscle (see Figure 6) shows the maximum tension that an individual muscle fiber can produce when the fiber is stretched to different lengths. This relationship is determined in the laboratory, using muscle fibers extracted from animals. The fiber is mounted at different lengths in a laboratory preparation, minimally activated to produce a tetanic contraction (the muscle contracts and *holds* its contraction), and the maximum tension produced by the muscle at a given length is recorded. As can be seen in Figure 6, the maximum tension that can be developed rises from low tensions at short lengths to a plateau at intermediate lengths, and with yet further lengthening the ability to generate tension progressively falls.




Figure 6. Skeletal length-tension relationship.

If you stretch any muscle excessively it's not going to be able to contract because you've pulled the actin and myosin completely apart, so that it can't form any cross-bridges and it can't contract, or, if you have a muscle that's too contracted to begin with, it can't contract anymore because the actin and myosin are pushed up against each other already and there's no room for them to slide any further.

Figure 7. Medical student discussing the length-tension relationship.

The misconception in component 1. The *misconception* held by students is that the differences in maximal tension potentially generated by an isolated skeletal muscle across its entire range of length (the length-tension relationship) are due to mechanical/anatomic factors, as follows. At lengths of fiber corresponding to the plateau, tension is greatest because of optimal alignment of actin and myosin filaments within the sarcomere of the muscle fiber. Force declines at longer lengths (descending limb of Figure 6) because, with stretching, potential binding sites are lost. Force declines at short lengths (the ascending limb) because of some kind of shortening-induced physical impediment to movement of the cross-bridges themselves, or to the physical structures that are moved during a contraction.

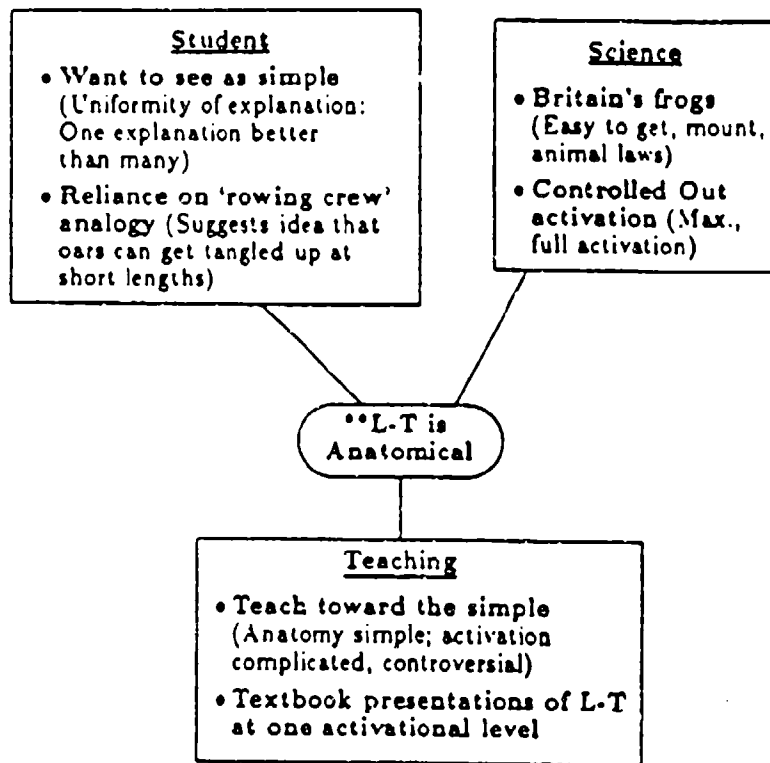
This supposition of an anatomical/mechanical basis for the entire range of the length-tension relationship is described in an example protocol from a medical student, given in Figure 7.

What's wrong? The characteristics of the length-tension relationship on its plateau and descending limb can be accounted for by the mechanically based Sliding Filament Theory, but this theory does not apply at all to the ascending limb; while the plateau and descending limb of the L-T are the product of optimal and (degrees of) suboptimal engagement of cross-bridges, the ascending limb (left-side) is not. The decline in tension by a muscle fiber at short lengths is a matter of diminution of muscle activation. In laboratory science, the activational basis for the ascending limb was revealed through the discovery of an automatic activational shut-down mechanism that engages when a fiber approaches short lengths when contracting; the very existence of short sarcomere lengths inhibits the activation. If means are used to override this shut-down mechanism, so that the fiber can be fully activated, tensions produced at fiber lengths corresponding to those on the ascending limb "spring up" to levels similar to those on the plateau (Jewell, 1977).

Hence, there are dual explanations underlying what is in appearance a fairly symmetric, continuous curve. The existence of this nonuniformity of explanation, within what appears to be such a clean, well-formed relationship, violates what we have identified as a prefigurative preference for uniform, integrative accounts.

Why Do Students Believe Component 1?

In discussing this and other component misconceptions, we will instantiate the framework for conceptual understanding (misunderstanding) developed in the first part of the chapter. Various parts of the framework will be instantiated, because the reasons for the misunderstandings are multifaceted. No single factor leads to misconception; rather, partial contributions, each reducing or simplifying the phenomenon in small ways, align in their simplification to yield robust and coherent, but erroneous, belief. These contributing factors involve the student, the educational process, and some practices of the biomedical science laboratory (see Figure 8).



****Implication:**

- Role of activation starts to get downplayed

Figure 8. Why do students believe Component 1? Converging influences.

Science. Before discussing the contribution of laboratory science to the students' misconception regarding the length-tension relationship, it is useful to mention the origins of this relationship itself. The classic length-tension relationship of muscle fiber that plays a key role in the misconception of heart failure is based on *skeletal* muscle from the hind-leg of *frogs*. It is interesting to note why this particular muscle was the one used in the late 1800's to establish this relationship. This muscle is large and easy to extract from frogs. Its structure is also conducive to easy mounting in the kinds of laboratory preparations required to determine tension produced at various lengths. Furthermore, the legislation in Britain responsible for the protection of animals in laboratory research at the time the experimentation was conducted did not (and still does not) consider amphibians to be "animals." Hence, for several reasons largely pertaining to *laboratory convenience*, the length-tension relationship of frog skeletal muscle exists as a classical teaching example for muscle function and shades students' ideas about how *heart* muscle works.

As we have noted, the tension developed by a muscle fiber is a function of both anatomy and activation level. Hence, activation is a third variable which modulates the tension developed at any length. In establishing the length-tension relationship within the scientific laboratory, activation was controlled to a single, maximum value. Hence, in this complexity reduction (under-dimensioning) of the response space of tension, the critical activational component is downplayed.

Students. Reductive tendencies on the part of the student also contribute to the idea that the basis for the entire length-tension relationship is mechanical. One is the desire for *uniformity in explanation*. This is reinforced by prescriptions about what good scientific explanation should be, for example, Occam's Razor, and maxims advocating singularity of explanation within medicine--"never two if one will do."

Another influence on the part of students is an over-reliance on a rowboat and rowing crew analogy (*reductive analogy*) for understanding dynamics of the muscle sarcomere (see the student protocol in Figure 3). This analogy conveys the idea that if the boat were shortened for a fixed number of rowers, the oars could get tangled up, or efficiency could otherwise be reduced by mechanical/anatomical obstructive factors.

Teaching. The mechanical components (related to anatomy) of tension production are easier to teach than the activational components; the mechanism of tension production by the cross-bridges (filaments) is less complicated. It involves only a few key elements (e.g., invariant filament lengths, shortening occurs by relative filament translation) and, especially when misinterpreted, is linear and, hence, easy to represent--the more potential binding sites, the more tension produced. In contrast, the process of activation involves numerous interacting subcellular components and processes (e.g., calcium ion release sites, binding proteins, ion pumps), and the effects are non-uniform (e.g., activation works qualitatively differently at different levels of anatomical extension). Second, as we have seen, neat, easily envisioned and cognitively productive analogies are readily available to aid anatomical understanding (for example, rowing crews and tug-of-war). Finally, details of the physiological activational account are less well understood and are more controversial in some particulars. Options and controversy are difficult to

deal with in both teaching and learning. (Anecdotes are widespread that medical students dislike competing or controversial accounts.) Hence, mechanical factors are likely to be stressed to a greater extent in teaching (reflecting, perhaps, yet another reductive bias: *Teach toward the simple*), better explicated, and almost surely more readily understood by the student.

Furthermore, textbook presentations of the length-tension relationship that students see and study depict the relationship at its single maximum activational level (*under-dimensioning*), reflecting the laboratory demonstrations discussed earlier. Three dimensional surfaces are difficult to represent in text.

Implications. The convergence of reductive factors such as those discussed leads to the de-emphasis of activational factors, in favor of mechanical/anatomical factors, in the production of tension (force) by the force-producing unit of a muscle. Playing a large role in this de-emphasis of activation is the overextension, in students' thinking, of the mechanical Sliding Filament Theory to the activationally based ascending limb of the length-tension relationship--thus eliminating altogether the need to consider activational aspects. But the muscle which has been discussed is skeletal muscle--not cardiac muscle--and it is from a fog! What, if anything, does any of this have to do with the heart?

Component 2: Cardiac Muscle is Like Skeletal Muscle: Both Have the Same Length-Tension Relationship (*Cardiac L-T = Skeletal L-T*)

The misconception in component 2 The functional dynamics of *cardiac* muscle fiber are assumed by students to be the same as for skeletal muscle (from the perspective of muscles of the body, an assumption of *homogeneity of components*). This includes the presumption that the two kinds of muscle generate the same kind of length-tension curve. Direct statements of this correspondence are generally not made by students (as it probably has the status of a tacit presupposition), but the correspondence is indirectly indicated by the description of the classical *skeletal* muscle length-tension relationship in discussions of the way the intact *heart* functions. (See the quote in Figure 9 and also the earlier textbook account in Figure 5 which more overtly equates the two L-T curves in its parenthetical statement.)

What's wrong? Heart muscle is very different from skeletal muscle. Two kinds of operational differences are most pertinent to the present discussion. First, heart muscle fiber produces a very different length-tension curve (see Figure 10, where the cardiac length-tension curve, solid line, is superimposed on the skeletal curve). In the cardiac curve, the potential for the development of tension continues to rise across fiber lengths where tension would fall in a skeletal muscle. In addition, there is *no corresponding cardiac descending limb*. In laboratory preparations used to establish the cardiac length-tension relationship, cardiac fibers are irreparably damaged or break

Previously I was discussing the length-tension relationship and bringing into focus the idea that it is the length of contact between actin and myosin fibers (sic: actin and myosin compose filaments within fibers) which up to a certain point as the length increases will have a corresponding increase in tension upon contraction of that muscle fiber. However, after a certain point the length of actin-myosin contact decreases because the muscle fiber is stretched beyond a certain point. As a result, tension after that point decreases because the muscle fiber is stretched beyond a certain point. As a result, tension after that point decreases as the fiber gets stretched and pulled past a certain length...This is responsible for the fact that as end-diastolic volume is increased Δh , and there is a corresponding increase in Δh , contractile force and stroke volume at systole.

Figure 9. Medical student discussing the basis for the Frank-Starling cardiac function relationship.

before they can be stretched to lengths corresponding to most of the skeletal descending limb. Second, cardiac and skeletal muscles activate and contract in very different ways. Skeletal muscles contract at maximal activation and produce tetanic contractions; they contract and hold the contraction until release. In contrast, cardiac muscles are always submaximally activated and they twitch, never achieving tetanic contractions.

Why Do Students Believe Component 2?

Science. Although cardiac and skeletal muscle are different in many ways, these differences were artificially minimized in establishing a length-tension relationship for cardiac muscle. Two means of reducing differences are particularly noteworthy (see Figure 11). First, the particular kind of muscle within the heart most similar anatomically to skeletal muscle was used for the demonstration. This is the papillary muscle which controls the action of heart valves. This muscle has a linear structure conducive to mounting in the laboratory and is amenable to a specification of length. It is not a muscle that pumps blood. In contrast, muscle involved in pumping blood, myocardial wall muscle, is *not* conducive to demonstrating a length-tension relationship. Such muscle is difficult to extract, difficult to mount, and is quite irregularly shaped, such that the concept of "length" makes little sense. Second, contrary to the normal submaximal twitch activation of heart muscle, pharmacological means were found to artificially "jack-up" the activation of the heart muscle fiber to produce the tetanic contractions of a length-tension relationship (Gibbs & Loisel, 1978). This further de-emphasizes the role of variable levels of activation in the production of force by the cardiac muscle. It also obscures the fact that because heart muscle is naturally submaximally activated, large ranges of force production are possible at *any* given (attainable) level of stretch.

Students. Although cardiac and skeletal muscle are different, they are also alike in some ways; for example, they both have sarcomeres containing actin and myosin filaments, force is generated by the binding and pulling of cross-bridges, and, within their respective dynamic ranges of length, more accessible binding sites yield more potential force. Given the similarity of skeletal and cardiac muscle on some of their features, there is an assumption by students of similarity on others (*extension of attributes*). In particular, given the similarities in the two kinds of muscle noted above, it is assumed they should be similar in their activational properties as well (which they are not). In its extreme form this kind of extension would lead to the idea that "all muscles are alike"; findings for any kind of muscle, whether skeletal, heart, papillary, or myocardial, are interchangeable (*homogeneity of components*).

Teaching. The classical *skeletal* muscle length-tension curve is the length-tension curve that is commonly used to introduce *cardiac* muscle dynamics in textbooks. Whether done for historical reasons (the skeletal L-T predates any cardiac one by 60 years) or for simplicity of explication, such use in instruction further encourages over-


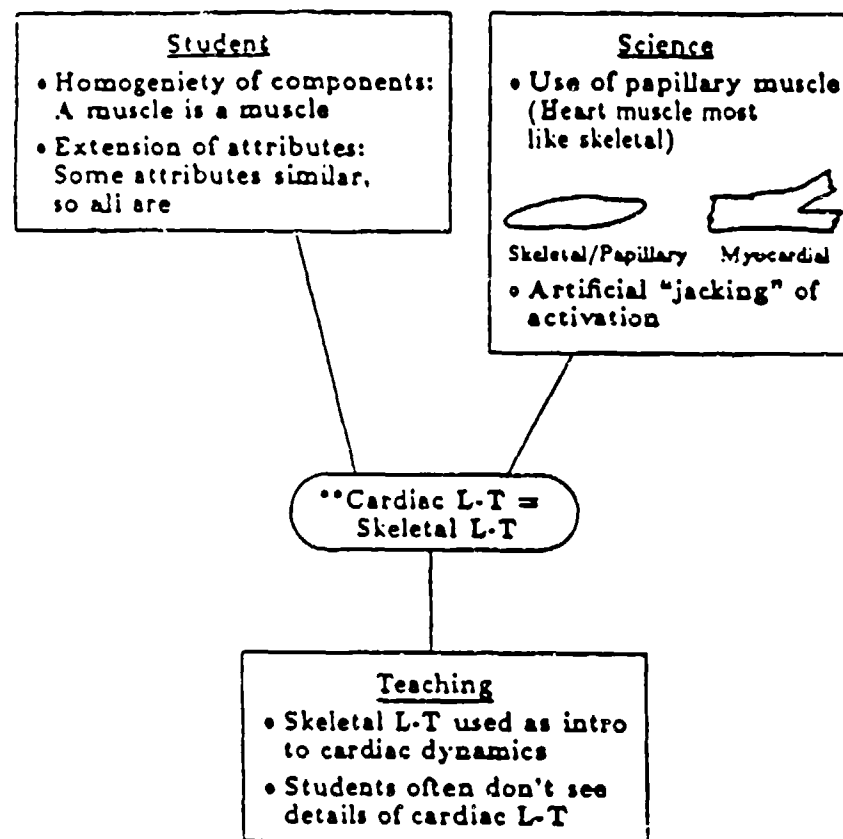


Figure 10. Cardiac length-tension relationship in comparison to skeletal length-tension relationship.



****Implications:**

- Believe heart muscle has a descending limb
- Twitch activation not in picture

Figure 11. Why do students believe Component 2? Converging influences.

attributions of similarity between the two kinds of muscle. Furthermore, textbook demonstrations of a *cardiac* length-tension relationship of any kind are rare. Hence, students do not have an opportunity to examine the differences and to incorporate the implications of the differences into their thinking about heart function.

Implications. Two consequences result from these multiple influences. First, students mistakenly believe that cardiac muscle, like skeletal muscle, has a descending limb on its length-tension curve, and hence, that cardiac muscle can lose force by being stretched to these ranges. Second, the activational factors that are critical to the production of force in cardiac muscle (but less so in operational skeletal muscle) are further de-emphasized. However, everything that has been discussed so far has been with regard to individual muscle fibers. To affect students' thinking about heart failure, they must be extended by students to the intact heart.

Component 3: Whole-Heart (Collective) Muscle Function is the Same as Individual Fiber Function; The Frank-Starling Relationship is a Notational Variant of the Length-Tension Relationship (*Frank-Starling* = *Length-Tension*)

Background to component 3. If an intact ventricle of the heart is filled with progressively greater volumes of blood, it produces a progressively more forceful contraction. This relationship, given in Figure 12, is called the *Frank-Starling relationship*, after the two individuals who simultaneously established it. The Frank-Starling relationship (F-S) as presented in Figure 12 is, again, determined by laboratory procedures. The curve is sometimes presented in textbooks with a small descending limb at large volumes (e.g., Katz, 1977, p. 204), but with no explanation of why force declines at these laboratory-induced, forced volumes. (Force declines because the heart muscle gets destroyed.) A caveat may (but often may *not*) follow that such volumes do not occur in the *in vivo* heart (Katz, 1977, p. 205).

The misconception in component 3. The *misconception* is that the Frank-Starling relationship for an intact ventricle is a direct reflection of the length-tension curve operating on the individual muscle fibers which make up the intact heart. By this account, the increasing force that occurs with increasing volume, on the ascending limb and "plateau" of the Frank-Starling relationship, occurs because component individual fibers are stretched to lengths on the ascending limb and plateau of the length-tension relationship. Force production declines at large volumes in the intact heart because individual fibers are stretched to the descending limb of their length-tension curves, as discussed in the protocol from a medical student given in Figure 13.

What's wrong? The direct mapping assumed by students, between volumes on the Frank-Starling curve and lengths on the length-tension curve, is erroneous. Part of




Figure 12. Frank-Starling relationship.

Starling's law states that as end-diastolic volume increases, the corresponding cardiac output at systole also increases. This occurs up to a point where the length of the fibers is too long and, therefore, tension at systole decreases causing a decrease in cardiac output. The decrease in contractile force corresponds with the length of myosin fibers being in contact with actin...as the length of contact between the fibers comes to a point where it starts to decrease because the cardiac muscle fiber is pulled to the length which is too long. That situation corresponds with the decrease in tension after a certain point is reached in the cardiac function curve.

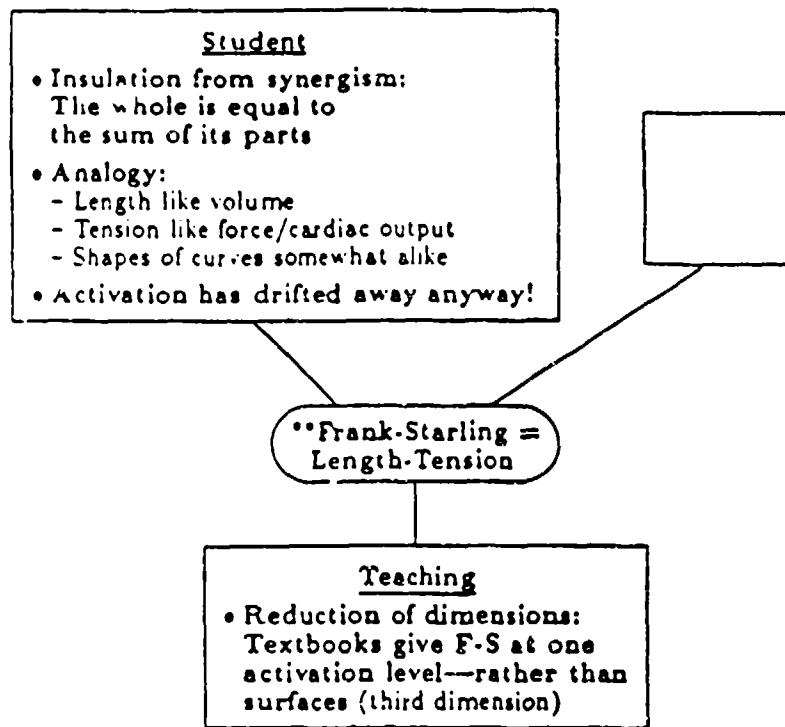
Figure 13. Medical student discussing the Frank-Starling relationship.

this mismatch is due to the geometric relation between length (or radius) and volume; large volume changes can occur with minimal changes in length. Hence, the dynamic range of lengths of fibers in an intact heart is greatly restricted, compared to the lengths occurring across the range of the laboratory-determined length-tension curve. In fact, most of the phenomenon of the Frank-Starling relationship occurs at fiber lengths corresponding to the plateau of the length-tension curve. Furthermore, the *length-tension relationship* that students (and others) assume in this misconception is the classical *skeletal* one (see Component 2) for, as we have seen, the corresponding cardiac length-tension curve has no descending limb. The mechanisms by which force is *actually* increased with increasing volumes in the intact heart (the Frank-Starling relationship) are, again, activational.

Why Do Students Believe Component 3?

Students. Of the many learner factors that might contribute to this articular belief (Coulson, Feltovich & Spiro, 1986), two seem particularly germane (see Figure 14). First, the students are assuming that the parts of a system in some sense "add-up" in combination to account for the function of the *intact* system (the heart): The whole is equal to the sum of the parts (*insulation from synergistic effects*). This is an example of what we earlier referred to as an analytic *prefigurative* reduction. Emergent properties that exist because of the structure and collectivity of the system (e.g., restrictions in the dynamic ranges of length, but also other effects such as hypertrophic adaptation--to be discussed later) are lost in this kind of analytic decomposition and attempt at additive reassembly. Second, even though the Frank-Starling relationship is not a collective counterpart of the length-tension relationship, there are enough enticing apparent similarities to lead students to treat them as at least analogical. Volume is *like* length in some respects (i.e., they are dimensions of *size*). Force (of a ventricle) is like tension. The respective curves for length-tension and Frank-Starling, showing the main relationships, look somewhat alike (especially if someone considers the *cardiac* length-tension relationship). Hence, it is easy to think of the Frank-Starling phenomenon as analogous to the length-tension phenomenon. The danger is that the detailed differences can be lost in the superficial analogy (*analogy treated as isomorphism*), as appears to be happening with the students. This is reinforced to the extent that activational contributions to force-production (the main basis for the Frank-Starling relationship) are neglected in favor of mechanical accounts, which, as we have argued, is encouraged from many sources.

Teaching. Textbooks typically present the Frank-Starling relationship as a two-dimensional curve, with the activational component implicitly held at one value (underdimensioning), rather than the three-dimensional surface that would be required if the activational dimension were to be represented as a variable (e.g., Katz, 1977, p. 204). This common simplification is, perhaps, a concession to the two-dimensional medium, but it serves to further de-emphasize the activational component. In addition, textbook accounts of the Frank-Starling relationship that do not make clear what is reflective of the laboratory and why (as noted earlier regarding the small down-turn in



****Implication:**

- The Frank-Starling relationship is believed to have an anatomical basis. In Frank-Starling, it is believed there are operative fiber lengths corresponding to the ascending, plateau, and descending limbs of the length-tension relationship.

Figure 14. Why do students believe Component 3? Converging influences.

the curve), versus what is representative of the *in vivo* heart, may also add to the problem (a form of *sanitizing*).

Implications. The implication for student thinking is that the Frank-Starling relationship for the intact ventricle is taken to be anatomically and mechanically based, in the same way that the length-tension relationship for individual muscle is taken to be. It is also mistakenly presumed that individual fibers in the functional heart can operate over the full range of lengths involved in the length-tension relationship for the sarcomere, including those corresponding to the ascending limb, plateau and descending limb.

Component 4: Heart Failure Results from the Overstretching of Muscle Fiber Filaments (*Heart Failure Results from Stretch*)

The misconception in component 4. Heart failure is thought to result from the heart becoming enlarged, stretching individual muscle units to lengths corresponding to those on the descending limb of the length-tension curve, and thus reducing the contractile force the ventricle can generate, as illustrated in the protocol of a medical student, given in Figure 15.

What's wrong? Many of the problems with this explanation have already been discussed, including the lack of a descending limb when the appropriate *cardiac* length-tension relationship is considered. Two other pertinent points will be addressed here. First, there are structural properties (Spotnitz & Sonnenblick, 1976) of the *intact* heart, including collagenous matrices, and also functional properties (Huntsman, Rondinone & Martyn, 1983) which prevent individual fibers from being stretched to lengths that would correspond to the descending limb of a length-tension curve. Second, an adaptive process of the heart, hypertrophy, intervenes when fibers approach overly long lengths, adding sarcomeres in series so that individual sarcomeres are returned to more normal lengths (Spotnitz & Sonnenblick, 1976). Both of these factors--the structural properties of fibers *in complex* and the hypertrophic adaptation process--exist at the level of the heart as an intact system, a level that is missed due to the analytic reduction that assumes that the system can be additively assembled from its parts.

Why Do students Believe Component 4?

Students. The heart in failure does enlarge, a classic sign of heart failure (see Figure 16). The question is whether it fails because it gets large, or, more appropriately, whether it gets large because it fails. This superficial "correlation" between heart and failure, we presume, encourages the inappropriate causal attribution, when the actual underlying causal mechanisms are not considered in detail (a reductive use of evidence: directional causality attributed to correlation). Furthermore, "bigness" (in the heart) and "long-length" (in a muscle fiber) share a common semantic dimension

Okay, the length-tension relationship of the muscle fiber...the length-tension curve. Is a curve that relates to the sarcomere of the muscle fiber; as tension increases your length is going to increase (sic, backwards: for the active relationship being discussed, an increase in length results in an increase in tension). There's a plateau and then it drops towards the end...the cardiac function curve or the Frank-Starling relationship is what I was just talking about (above), the cardiac function curve goes up and then plateaus out...I'm just thinking of a curve in my mind where I see there's a certain level beyond which when it drops down it puts you into congestive heart failure.

Figure 15. Medical student discussing heart failure.

that encourages and supports the attribution, perhaps through cognitive activation spread (Anderson, 1983; Collins & Loftus, 1975), whenever either the large heart of heart failure or the length-tension relationship of muscle is considered. Also, the properties of the intact heart that prevent overstretch (collective *structure* of the heart, and the collective *function* of hypertrophy) are not likely to be accounted for well in thinking about heart failure if students assume that whole heart function can be assembled from its constituent components.

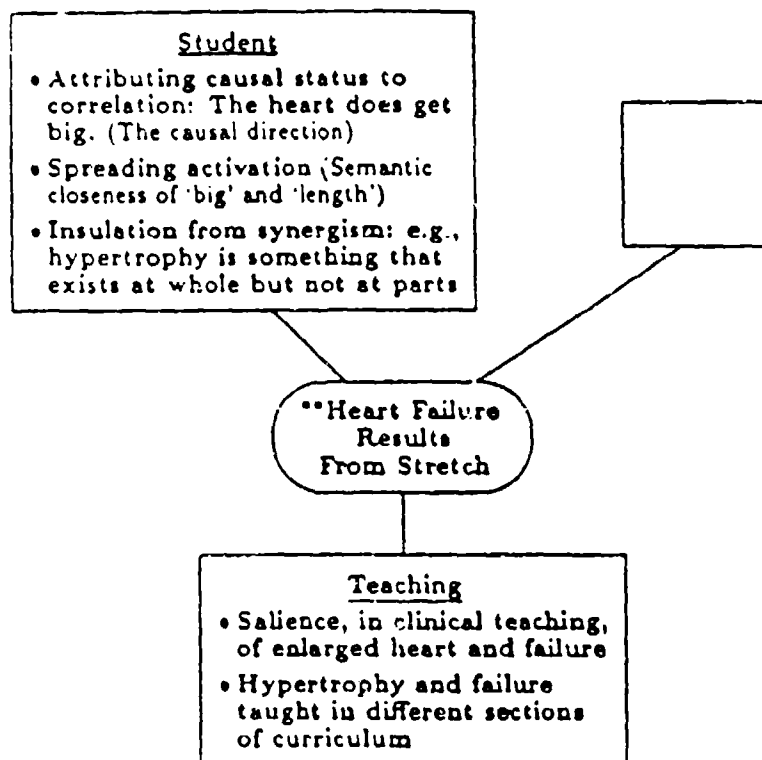
Teaching. The enlarged heart of heart failure is highly salient in clinical teaching about heart failure. It is one of the classic clinical signs monitored to determine the presence and severity of heart failure. The salience of this feature, interacting with other related features and interpretations we have mentioned, may serve to reinforce the idea of failure as resulting from "stretch." In addition, for convenience and clarity of exposition, the interrelated processes of hypertrophic adaptation and heart failure are often taught in different sections of texts (*compartmentalizing*).

Implications. The implication for students' thinking is that heart failure is believed to result from stretching individual myocardial fibers to lengths at which, for anatomical/mechanical reasons, they cannot generate adequate force. The heart is mistakenly believed to fail because it "falls over the edge" of the Frank-Starling curve, onto a "descending limb" that is a direct reflection of the descending limb of the skeletal length-tension curve. Again, Sliding Filament Theory dynamics that account for the descending limb of the skeletal length-tension curve are often extended, this time to the functional heart.

Activational contributions that are in fact responsible are not sufficiently considered. This is *not* to say that students do not know anything about activational components in muscle function. When students are directed to activational dynamics in our studies of muscle activation and contraction, they often discuss aspects of these dynamics appropriately. The point here is that in considering *heart failure*, for many students these activational dynamics are not integrated into the process as they should be, and the inappropriate anatomical/mechanical account lingers in their thinking. It is worth noting that the lack of integration of anatomical and activational factors that was found in students' thinking about the length-tension relationship is recapitulated in their thinking about heart failure.

3.1.3.5 More than a Simple Chain of Reasoning? A Network

It is easy, in a way, to understand the seductive plausibility of the explanation of heart failure as resulting from stretch. Why does heart failure happen? The heart is made up of individual muscle fibers; there is a sense in which individual muscle fibers fail because of stretch; so why not the heart which is composed of fibers? Is this misconception simply the result of a faulty chain of reasoning of the following sort?



****Implication:**

- Heart failure results from stretch, has an anatomical basis

Figure 16. Why do students believe Component 4? Converging influences.

LENGTH-TENSION IS ANATOMICAL (STRETCH) CARDIAC L-T =
SKELETAL L-T FRANK-STARLING =
LENGTH-TENSION HEART FAILURE RESULTS FROM STRETCH

The skeletal length-tension relationship is anatomically (stretch) based; the cardiac L-T is like the skeletal one; the Frank-Starling relationship for the whole cardiac muscle is like the individual L-T; therefore, the whole heart fails because the individual fibers become stretched. A reasoning chain such as this probably accounts in part for the misconception. But, that is not all there is to it. As we have seen, there are holes in the chain (e.g., the existence of a descending limb for the cardiac length-tension curve) that seem to get filled in from elsewhere to be consistent with the overall chain. In addition, why is the *overall* belief so powerful that it overrides contrary explanations from the students' main curricular textbook?

Our hypothesis is that bits and pieces of knowledge, in themselves sometimes partly correct, sometimes partly wrong in aspects, or sometimes absent in critical places, interact with each other to create large-scale and robust misconceptions. Such a structure of knowledge could be represented as an interactive network, where fragments of knowledge connect in complicated ways to mutually strengthen or weaken each other and to produce aggregate knowledge that is functionally different (in this case stronger) than the pieces (McClelland, Rumelhart & the PDP Research Group, 1986; Rumelhart, McClelland & the PDP Research Group, 1986; Waltz, 1985;).

A depiction of this kind of knowledge structure, representing aspects of the heart failure misconception, is given in Figure 17. Interactive and synergistic effects occur in this network in several ways, as illustrated in the following examples. One form of interactive effect involves *extension of attributes* among entities, as has already been discussed in this chapter. If two entities have many attributes in common, then the likelihood of their similarity on others is enhanced. For example, if skeletal and cardiac muscle are similar in many respects, then it is easier to believe inappropriately that they both have similar activational properties, that they both have a descending limb on their length-tension curves, and so on. A second interactive effect involves *reciprocation*. Belief in A makes it easier to believe in B, and *vice versa*. Belief that cardiac muscle has a descending limb on its length-tension curve makes it easier to believe that an *in vivo* ventricle would lose force for mechanical/anatomical reasons at large blood volumes, and that this would be reflected in a downturn in the Frank-Starling relationship at large volumes. In turn, belief that the Frank-Starling relationship has a downturn in force at large blood volumes that are naturally possible (a belief abetted in part by ambiguous textbook treatments of this matter, as we have noted) bolsters belief in a (nonexistent) descending limb for the *cardiac* length-tension curve (especially if one knows about the *skeletal* L-T relationship).

Another network-wide effect involves *pervasive colorations*. Colorative effects occur because of the presence in many places within the network of aspects of knowledge that bear a "family resemblance" to each other. An example from the heart failure

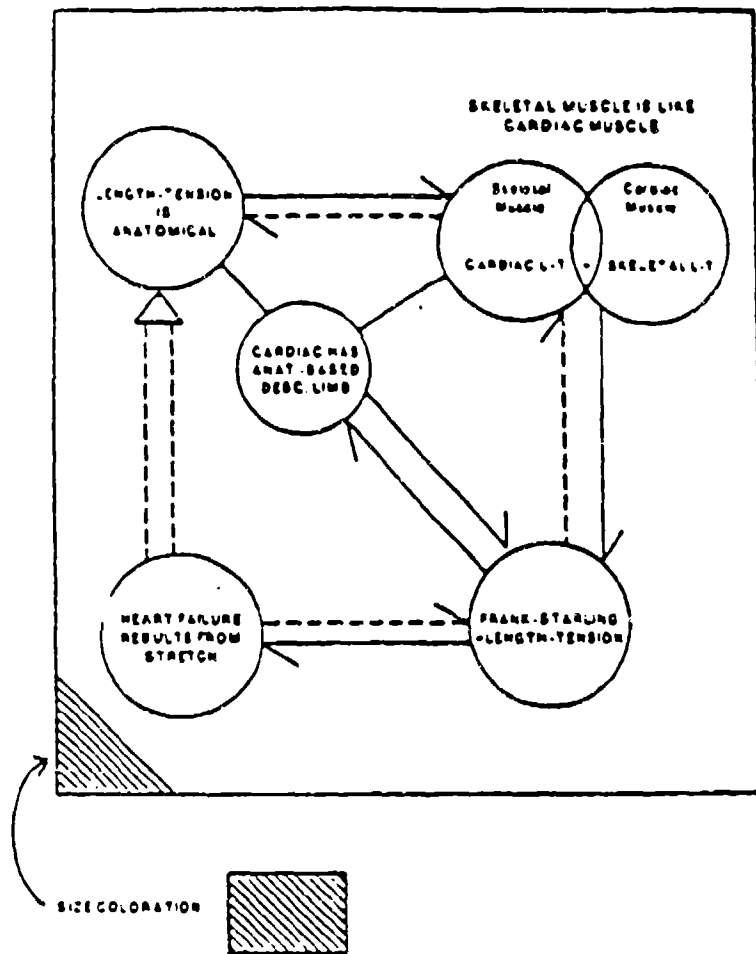


Figure 17. The misconception; a network of reciprocating beliefs.

misconception is the aspect of *size*. It is present in the length-tension relationship through *length*, in the Frank-Starling relationship through *volume*, and in the *bigness* of the failing heart. The presence of a common semantic dimension of size in all facets of muscle function reinforces both the importance of *size* generally and misconceptions associated with it within the facets themselves. As *size* emerges progressively (from many sources) as a recurrent theme during the course of the development of the misconception, it comes to take on the ability to bolster belief in the misconception independently of the 'reasoning chain' and to support all constituent misconceptions related to size.

Why is the reasoning chain portrayed earlier in this section so seductively plausible? The answer, we believe, raises yet another type of synergistic effect--that of *structural corroboration*. The argument is alluring because it is so efficiently compact and tidy. But, this "good form," is achieved only because the ideas contained within the reasoning chain have already been misconceived in ways that all align toward internal coherence across the network, enabling the "syllogistic" overlay. The pervasive colorative effect of size across the whole network inlay adds to this structural integrity. Structural corroboration refers to the contribution to belief that devolves from multifaceted well-formedness, accomplished by the widespread alignment of elements within the network.

Finally, yet another way that interactive effects are manifested is in the fact that reasoning within the network can proceed in *multiple directions*. We have seen how an initial focus on the length-tension relationship can lead to the inappropriate account of heart failure. But, what if a person starts with a focus on heart failure itself and tries to decipher a cause? It is easy, especially if the contribution of activation is not recognized appropriately, to step backward and to find a plausible account in the descending limb of the length-tension relationship.

The structure we propose for the inappropriate conception of heart failure (and, perhaps, for the nature of complex ideas in general) is a connected, interactive, reciprocating network. Fragments, pieces and partial dimensions of knowledge feedback on each other in nonlinear, reciprocal ways. In the case of the present misconception, everywhere the network is "jiggled" it broadcasts dysfunctional influences. The result is a whole that is stronger than the sum of its parts. Various *simplifications* of complex phenomena play a large part in the acquisition and maintenance of the (in *this case*) dysfunctional structure. These include simplifications on the part of the learner, perhaps to aid some form of coherent comprehension; on the part of the teaching process, perhaps to ease students into sophisticated understanding; and on the part of laboratory science, perhaps within the incremental quest for understanding of complex phenomena (but where consumers of research, e.g., students, see and are affected only by the partial products of the overall quest, without access to the "big picture").

3.1.4 Implications for Traditional Sequencing of Instruction

In this chapter we have discussed our findings concerning patterns of pervasive failure in the development of advanced conceptual understanding, with an emphasis on

the contribution of oversimplification to the development of error. These findings have implications for traditional approaches to learning and instruction.

When the curriculum is dense, the ideas are difficult, and the pace is fast, all of this promotes attempts at simplification--on the part of both learners and teachers who must cope with the pace. In simplifying, students can gain some level of coherent and satisfying understanding, teachers can "get through" the material, easily scorable tests can be built and graded, and so forth. There is probably a faith that simpler initial understandings can be built on progressively through the curriculum. In some instances such incremental approaches to instruction can be effective. However, under certain conditions, such approaches are susceptible to hazards of the kind we have shown in this chapter.

-- First, oversimplified initial versions of a concept can produce a false sense of understanding and abort the pursuit of deeper understanding. When a concept is especially complex and multifaceted, this problem can become more acute, because partial misunderstandings can reinforce each other.

-- Second, instructional efforts to challenge and change a student's oversimplified conception (to raise it to a higher level of sophistication) may fail. The student will minimize discrepancies with the simpler, cognitively satisfying model so that this model can be retained. The student may not notice discrepant aspects of the concept or, if they are noticed, they may be filtered in interpretation toward the model that is already held. Local adjustments in this model will take precedence over fundamental reformulations.

-- Third, educational strategies that attempt to teach complex concepts by focusing on their simpler components (with the hope of building toward fuller understanding) may encounter an additional kind of problem. Some concepts (and their components) are inherently inextricable from their organic functional context (as is the case in the phenomenon of heart failure). Concepts such as these involve, for example, synergistic properties or interactions among numerous variables. For concepts of this kind, any analytic decomposition misrepresents the concept fundamentally. One cannot make the components "add-up" to the whole, and there is no alternative in instruction but to find ways to convey the irreducible complexity in a manner that is tractable.

-- Fourth, as we have seen, incremental approaches that start with simplifications can engender associated habits of thought and learning that remain and interfere with advanced knowledge acquisition. An important focus of the research described in Section 4 of this Report is on finding ways to convey complex material tractably, without having to resort to oversimplification, sometimes even at the earliest stages of instruction.

3.1.5 Footnotes to Chapter

Paul Feltoovich, Rand Spiro, and Richard Coulson were the main contributors to this chapter.

¹We thank Alan Lesgold for discussing this issue with us.

²Our method of studying students' understanding of a concept has two major components. One is a scheme for analyzing conceptual structure, which can be used to identify areas of potential cognitive difficulty. This scheme has both analytic and synthetic components. Because contributions to misunderstanding of a concept may occur because of errors in understanding of even its most basic elements, the analytic part involves breaking the concept into its primitive elements, typically concrete entities with physically realizable properties (e.g., definitions of "myosin" and "cross-bridges" with regard to cardiac muscle dynamics). The more synthetic aspect involves the combination of these elements according to higher-order relationships among them, and emergent, more abstract conceptual aspects not easily tied directly to the basic "building blocks."

For each concept studied, a probe set of discussion questions is created for use with students in the laboratory. These probe sets for a concept are closely tied to the conceptual analysis, and all probe sets have a similar general form. This form can be thought of as an hour glass, which starts general and open-ended, comes down to specific questions, definitional elements, etc., and then builds back up again to conceptual component combinations and clinical applications.

In particular, the first probe question is always the same and addresses analogies and other kinds of models a student may use in thinking about a concept. The second question is a full and open-ended discussion question that spans the entire concept of focus. This is included early in order to gain an appraisal of the student's understanding before any aid or prompting that may result from later items of the probe set. The discussion questions then come down to basic elements of a concept. The questions then expand into progressively higher aspects of the concept under focus. The final items of a probe set include carefully selected application questions, often including questions chosen to reveal classes of misconception that can be envisioned *a priori*.

These probe sets are used with medical students (and practitioners) in individual sessions in the laboratory. Students discuss each question, in order, for as long a time as they wish. This first pass through the set is observed by at least one project member, and there follows a period of directed questioning by the observer prompted from the student's discussions. At the conclusion of the session there follows a period of open discussion of the concept. A session for one concept (and student) usually lasts about 2 hours. The entire session is audio-taped for transcription.

Analyses are directed at various kinds of commonalties in responses, at patterns of interdependency among responses to probe set items that suggest coherent conceptual models, and at patterns of response across probe sets for a family of related concepts that serve to corroborate (or disconfirm) our ideas about a student's conceptual model or its aspects.

³We thank Dedre Gentner for suggesting this term for the phenomenon.

⁴This reduction will be recognized by the reader as resembling a cognitive operation that is commonly portrayed as useful, enabling prediction of attributes in the

1

absence of specific evidence, or more generally contributing to cognitive efficiency ("default values" in "frames" or "schema" theory, e.g., Minsky, 1975; "default assignment" of missing values of an object based on its shared values with other objects in more recent PDP models, e.g., McClelland, Rumelhart & Hinton, 1986). The status of such operations as being common to people and as sometimes being serviceable *because* they enable a kind of efficiency is not at odds with our more negative treatment of them here. Our argument is that such generally efficient effects, while perhaps having a serviceable function in many simple and routine knowledge domains, may actually be a major hazard or *impediment* when it comes to learning difficult, complex, sometimes abstract, sometimes *counter-intuitive* ideas of the sort we are discussing. Such ideas are hard to understand partly *because* our mundane cognitive apparatus, honed in (one might even say "designed for") the commonplace, is in many ways at odds with what is needed. One might argue that the more that ideas are "out of synch" with common cognitive mechanisms, the more difficult they will be.

*While such a treatment is not appropriate for this chapter, the reasons why this mechanical "overstretching" account of heart failure *cannot hold* are explicated in great detail in another chapter (Coulson, Feltovich & Spiro, 1986).

*Students participated shortly after having had the most pertinent (cardiovascular) part of their curriculum in each year. These proportions are conservative since only clear statements of the misconception were counted.

Chapter 3.2 Analogy-Induced Misconception in Advanced Knowledge Acquisition: A Detailed Study of One Kind of Maladaptive Reductive Bias

In the last chapter we examined the range of reductive biases that lead to misconception development and the larger patterns of faulty learning that result. In the present chapter we adjust our grain size, studying in detail the learning deficiencies that result from just *one* kind of reductive tendency in the learning of complex material: the overreliance on analogy to prior knowledge.

3.2.1 Overview of Chapter and Introductory Remarks

Overview

A pervasive tendency for analogies to contribute to the development of entrenched misconceptions is identified. The misconceptions have the form of reducing complex new knowledge to the core of a source analogy. A taxonomy of ways that analogies induce conceptual error is presented, with examples of common biomedical misconceptions found in our research corresponding to each. The factors in learning and instruction that contribute to the misuse of analogy in advanced learning are examined. Finally, the situation of analogy is argued to be one instance of a more general pattern of oversimplification-based misconception development attributable to the use of single knowledge sources when compilation of multiple sources would be more appropriate.

Introduction

Few would disagree that analogy is an important tool in the acquisition of new knowledge. Indeed, work in cognitive science and educational psychology in the last dozen years provides ample evidence of the usefulness of analogy in learning and has substantially advanced our understanding of the psychological mechanisms responsible for that utility (e.g., Burstein, 1983; Carbonell, 1983; Collins & Gentner, in press; Gentner, 1983; Gentner & Gentner, 1983; Gick & Holyoak, 1980; Rumelhart & Norman, 1981; Vosniadou & Ortony, 1983). Yet, as this chapter will demonstrate, the use of analogies in learning is far from straightforward and surprisingly often results in deeply held erroneous knowledge.

Our intention is to offer a more temporized and cautionary alternative to the general enthusiasm for learning by analogy, especially in its most common form: The use of a single mapping between a source and a target concept (the "topic")--what we will refer to as a "single analogy." (For exceptions that address more complex uses of analogy, see Burstein, 1983; Collins & Gentner, in press.) We argue that single analogies that help novices to gain a preliminary grasp of difficult, complex concepts may later become serious impediments to fuller and more correct understandings. Specifically, although single analogies rarely if ever form the basis for a full understanding of a newly encountered concept, there is nevertheless a powerful tendency for learners to continue to limit their understanding to just those aspects of the new concept covered by mapping it to the old one. Analogies seduce learners into reducing complex concepts to a simpler and more familiar analogical core.

Our position is not antagonistic to analogy; again, there is no doubting the value of apt analogies in introducing unfamiliar concepts. However, we are not as sanguine about the benefits of single analogies at later, more advanced stages of learning about complex concepts. Therefore, sensing an unhealthy imbalance in the attention devoted to the nature and benefits of analogical learning, we attempt to address some of the more neglected hazards. On a more positive note, we suggest two antidotes for what we have found to be an insidious tendency of analogies to block more advanced knowledge acquisition: (a) pay more attention to the ways that analogies fail or mislead or are incomplete--learners and teachers are more likely to be able to avoid pitfalls if they have explicit warning of what those pitfalls are, and (b) employ integrated multiple analogies (see Chapter 4.2) to convey more of the necessary complexity of difficult concepts--the more complex and ill-structured the new concept, the greater the need for a finely tuned synthesis of the relations between it and *several* already known concepts.

Plan for the Chapter

We begin with a brief introduction to the overarching orientation that guides our remarks on analogy. Then we illustrate the danger of misuse of single analogies in the learning of complex concepts, using examples drawn from the biomedical domain. We demonstrate several common misconceptions held by medical students that are traceable to a cognitive (and sometimes instructional) overreliance on single analogies. Eight varieties of analogy-induced misconception are identified. We then examine the circumstances of learning and instruction that promote the uncritical acceptance and entrenchment of learning based on inadequate analogies.

3.2.2 The Perspective of Advanced Knowledge Acquisition

The work on analogy that we discuss in this chapter is part of the larger program of research concerned with advanced knowledge acquisition discussed throughout this Report. Here we briefly review some of the general issues. Advanced knowledge acquisition refers to the learning that follows initiation into the rudiments of a knowledge domain and precedes the attainment of expertise. This intermediate stage, falling between the novice and the expert, is often neglected. We argue that this neglect has serious consequences because the aims of advanced knowledge acquisition are different from those of introductory learning, and those differing aims are best attained by qualitatively opposed means--what helps at the introductory stage may hurt at the advanced stage, and vice versa. In other words, success at the introductory stage may sometimes result in forms of entrenched fundamental understanding that interfere with the eventual attainment of expertise.

The main aims of advanced knowledge acquisition are:

1. mastery of complexity: the acquisition of those aspects of conceptual complexity that are necessary for a correct understanding of important concepts (rather than the attainment of a superficial familiarity with simplified versions of concepts, which is often the goal at introductory stages of learning in a knowledge domain); and, relatedly,
2. knowledge applicability: development of the ability to adaptively apply or transfer acquired knowledge, especially to realistic situations, especially as those situations

present aspects of novelty (rather than merely being able to reproduce content material from memory in just the way that it was learned, the criterion too often employed at all stages of learning).

The present chapter is concerned with the first aim. We report on misconceptions held by advanced learners (medical students) that all involve a failure to master important complexities in concepts (because of interference from an analogy); and we present an approach to mental representation, learning, and instruction that is designed expressly to address the difficulties posed by complexity. (Although full apprehension of conceptual complexities is often a necessary but not sufficient condition of knowledge application--the second aim--that issue is not addressed in this chapter.) Of course, not all concepts are important enough that their complexities must be mastered. However, at some stage in knowledge acquisition a point is reached where certain concepts are so central that they must be correctly understood. Furthermore, we argue that the complexities of a knowledge domain become much more centrally important to the extent that knowledge must be applied in unconstrained, naturally occurring situations (rather than learned in the abstract or for application in artificial instructional settings). At some point during the progress of learning in a content area, it becomes important that the learner "get it right," even if the resulting difficulties place burdensome demands on learners and teachers.

To insure our not being misunderstood, it is worth repeating at this point that we are not advocates of complication for its own sake. Well-structured subject matter can and should be taught in much simpler ways than we recommend for more complex and ill-structured material; every kind of reductive use of analogy that we criticize has a counterpart which is a strategy for usefully stripping away complexities--either at the early stages of learning some topic (especially for learners who will not be going further in a subject area), or when complicating factors are not especially important. Furthermore, we realize that the premature introduction of complexity may confuse learners. However, we do argue that complexity should be introduced as soon as learners are ready for it, because early simple frameworks often act as impediments to grasping complexities introduced later (as we shall see in the next section). Furthermore, new approaches to learning and instruction are now being developed that will advance the time at which learners are ready to deal with complexity--procedures are being created for making the learning and instruction of complex information more tractable at an earlier stage of knowledge acquisition. One example of such a procedure is the use of integrated multiple analogies, made psychologically manageable with composite images; this procedure is dealt with in Section 4 of the Report.

3.2.3 The Reductive Force of Analogy and Analogy-Induced Misconceptions

Those who study analogy offer frequent reminders that the useful applicability of an analogy is never total; only some relational aspects of a source domain may be transported to a target domain (the topic). For example, only a subset of the relationships in the domain "solar system" are representative of the domain of "atoms" (Gentner & Gentner, 1983). While this point is obvious to those who study analogy, it unfortunately fails to characterize the state of affairs with respect to the actual employment of analogies in learning and instruction. In fact, as we shall see, there is evidence that analogies exert a powerful "reductive force": When a striking, pedagogically efficient analogy is employed that incompletely represents some target

of understanding, the incomplete representation often remains as the only representation of the target concept. We have found that misconceptions attributable to the reductive effect of analogies occur even when teachers and texts are explicit in stressing the inadequacy of an analogy. In other words, when analogies are used to "start simple," the knowledge ultimately acquired often stays simple. Well intended analogies often result in oversimplified knowledge.

3.2.3.1 Analogy-Induced Misconception: A Typology of Varieties, with Biomedical Examples from our Studies

We have used directed, open-ended discussion probes to assess medical students' knowledge. (The probe procedure is discussed in Coulson et al., 1986, and in Chapter 3.1.) The probes have uncovered several instances of commonly held misconceptions connected to cognitive and instructional aspects of the use of analogies. We will simply refer to these misconceptions without presenting methodological details and specific results of the various studies that they have been culled from (i.e., the examples of analogy-induced misconception have been taken from research in which issues related to analogy happened to emerge in the context of probes with a more open-ended focus or a focus directed toward matters other than analogy). More detailed treatments of the misconceptions discussed here and the procedures in the studies that revealed them are available elsewhere (Coulson et al., 1986; Fellovich et al., 1989; Spiro et al., 1988).

Eight ways that analogies can induce misconceptions are identified below. Each type is illustrated by an example of a biomedical misconception that we have observed to be common among medical students (and, occasionally, among physicians and popular medical textbooks). Since the analogy-based misconceptions involve very technical subject matter, our characterizations will stress the form or structure of the fallacious knowledge and how it develops, rather than going into full detail on the specific content of the misconception. (Again, references in which more substantive discussions are provided for each misconception have been provided.)

All of the examples have two features in common: (a) the source (or base) domain information in the analogy is inadequate or potentially misleading for understanding the target domain (the topic), and (b) in practice, the knowledge acquired about the topic is reduced to just that information mapped by (inadequate) analogy from the source domain. This includes both incorrect overextensions from the source (derived from misleading aspects) and omissions in the source of information important for understanding the topic.

By presenting this list of types of analogy-induced misconception our dual intentions are to call attention to such effects of analogy and, as a step toward prevention, to provide a somewhat detailed analysis of the forms in which these deleterious effects are manifest, thereby making learners and teachers more alert to their occurrence. It is not claimed that the eight types identified form an exhaustive list or that some misconceptions may not be characterizable according to more than one type. Also, the order in which the types are presented is dictated primarily by sequencing requirements related to the biomedical examples, rather than by any natural ordering of the misconception types.

1. Indirectly misleading properties. Some salient characteristic of the source domain that is not central to the pedagogical point of the analogy adversely influences understanding of a parallel characteristic in the topic domain. Roughly put in schema-theoretic terms, there is a prominent variable "slot" in the source domain that has a different instantiation in the topic domain (Rumelhart, 1980). As will become clear after the following example, it is not the case that the slot's value is changed in the topic domain to be identical to its value in the source; rather, misconceptions develop regarding the topic that derive from properties entailed by the mismatched slot instantiation in the source.

EXAMPLE: A common analogy used to teach opposition to blood flow (impedance) uses rigid pipe systems such as household plumbing as the source domain. This analogy promotes understanding of that aspect of impedance due to resistance, which primarily depends on the radius of the vessel. However, unlike plumbing pipes, blood vessels are flexible. Related to that flexibility is an additional source of impedance, compliant reactance. Further, a third source of impedance exists, inertial reactance, which also derives from the pulsatile (beating) nature of the heart, accelerating the mass of blood on every beat. These latter two forms of impedance are jointly referred to as reactance. Aspects of impedance that involve reactances are frequently misunderstood by medical students (Feltovich et al., 1986, in press). The misunderstandings take the form of either ignoring reactance, or (mistakenly) reinterpreting reactance phenomena in terms of their limited aspects that bear resemblance to resistance (where, again, resistance is the only aspect of impedance supported by the plumbing analogy). For example, compliance (the stretchiness of vessels--related to compliant reactance) is erroneously thought to contribute to impedance through the ability of a stretchy vessel to change radius and thereby affect resistance; and blood density, which directly affects inertance (and inertial reactance), is either neglected by students or it is equated with blood viscosity, which is, again, a contributor to resistance. This family of misconceptions connected to the rigid pipes analogy occurs despite the fact that students usually are exposed to a complete account of the factors contributing to impedance.

It is very important to note that the students do not have the misconception that blood vessels are rigid. The misconceptions that we have referred to come from those aspects of the analogy to household plumbing pipes that are entailed by their rigidity. So, for example, the rigidity of plumbing pipes (as well as the constant, as opposed to pulsatile, pressure head that is usually associated with plumbing) make resistance the only factor substantially opposing flow, and misconceptions about impedance to blood flow tend to involve the erroneous conversion of nonresistance phenomena to ones that are resistance-like. Thus the effects of the misleading slot instantiation in the source (rigidity) on the development of misconceptions of the topic are indirect. Of course, there may be cases where the learner mistakenly adopts for the topic domain the actual value of a mismatching variable slot from the source domain. This would be more likely to happen when the instantiation of the variable slot is not especially important or salient (in contrast to the obvious actual flexibility associated with blood vessels).

2. Missing properties. An important aspect or characteristic of the topic domain has no counterpart in the source domain, and that missing aspect or characteristic does not get incorporated in the understanding of the topic.

EXAMPLE: Medical students frequently have trouble attaining a sophisticated understanding of pressure in the cardiovascular system. A contributing factor is, again, the analogy to household plumbing, water taps, and so on. Because these more familiar fluid systems have a constant pressure head, there is no need in these systems to think about the variable acceleration of water. When the plumbing analogy is used in the heart domain, this missing aspect leads students to omit from their thinking contributions to blood pressure deriving from the pulsatile acceleration of blood. The conception of pressure and the factors that influence it are thereby reduced to ones that would be captured by the familiar Ohm's law, applicable in plumbing.

3. Exportation of base domain properties. A salient characteristic of the source domain that has no analogue in the topic domain is nevertheless exported to the topic. A non-existent "slot" is created in the topic to correspond to a slot in the source.

EXAMPLE: Starling's relationship between end diastolic volume in the heart and cardiac output is seen as analogous to the relationship between the length to which an individual muscle fiber is stretched and the tension it can produce. This is partly because the graphs of both relationships (volume-output and length-tension) have similar ascending left limbs and plateaus. However, while the length-tension relationship for an individual (skeletal) fiber has a descending limb reflecting decreased tension at long lengths, the Starling relationship (in the in vivo heart) has no corresponding descending limb at large volumes. A common error made by medical students, physicians, and some medical texts is to assume that the Starling relationship has a descending right limb like that of the length-tension curve for skeletal muscle (Coulson et al., 1986).

This mistaken importation of the source domain's descending limb has serious consequences, playing a central role in the development of a major misconception about the nature of congestive heart failure. Heart failure is erroneously attributed to "falling off the plateau of the Starling curve" as the collection of individual muscle fibers get stretched too far (thus losing tension), resulting in reduced cardiac output, further stretching because of blood accumulation in the heart, further reduced output, and so on. In fact, there are physiological limitations preventing the cardiac muscle fibers from reaching a "descending limb" on the Starling curve (for an in vivo heart). In contrast to this misconception, the heart enlarges due to hypertrophic mechanisms while it weakens as a result of the breakdown of muscle activation mechanisms--the heart does not fail because it is getting enlarged (stretched too far), as the analogy to the individual skeletal muscle fiber length-tension relationship mistakenly suggests; it gets enlarged in response to its failing. (See Coulson et al., 1986, for further explication.)

(As we will see in the next section of the chapter, "The Analogical Bias," "length-tension: Starling's volume-output" is a false analogy--the circumstances of its mistaken adoption will be discussed later. Our present point only examines its consequences after it has been mistakenly adopted.)

4. Directly misleading properties. A nonsalient aspect of the source domain in an analogy has a different value than the parallel aspect of the topic domain. The variable slot in the topic is incorrectly assigned the instantiated value of the slot in the source.

EXAMPLE: Recall the misconception that heart failure is due to collective (ventricular) overstretching of individual muscle fibers, "falling off the plateau of the Starling curve," in much the same way that such "falling off" could be achieved in an individual skeletal muscle fiber's length-tension curve (Coulson et al., 1986). The correct explanation involves biochemical activation or energizing of the components of the system. One reason that activation is neglected by medical students as a possible cause of heart failure is that cardiac muscle fibers are understood by analogy to skeletal muscle fibers. The implicit assumption is that "all striated muscles are alike." In reality, the two kinds of muscle differ greatly in their activational properties, in ways that are important with regard to heart failure. Skeletal muscle always functions at full activation, so that the activation level of the muscle is not an issue. If cardiac muscle is taken to be like skeletal muscle, and the level of activation is not an issue for skeletal muscle, then it follows that activation is not an important issue for cardiac muscle either. Unfortunately, in contrast to skeletal muscle, the degree of activation of cardiac muscle is variable. And degree of activation is the key to the correct account of heart failure (in a way that would not be similarly manifested in skeletal muscle failure). The analogy of cardiac to skeletal muscle brings along with it an incorrect instantiation of the variable slot, "type of activation," and adopting that property of the source domain for the topic has serious ramifications.

5. Focus on surface descriptive aspects with corresponding mistreatment of underlying causation. Some analogies are very effective at characterizing surface features and relationships, but gloss over underlying causal mechanisms. The result is that learners are susceptible to either filling in a convenient but incorrect causal account of their own, or just leaving the causal mechanism unexplained, as a kind of "black box." (It might be said that a comparison based primarily on surface descriptive aspects is more metaphorical than analogical. However, our point here is that an underlying relational structure is indeed transferred--that is, people have a tendency to interpret metaphor analogically.)

EXAMPLE: The failing heart is often compared by medical students to a stretched out, saggy balloon (or to overstretched "silly putty"). In fact, this analogy gives a very good picture of representational correlates in the two domains; both heart and balloon get big, but the heart gains passive tension while the balloon loses it. The stretched out balloon, when underinflated for its enlarged size, exhibits floppy walls (low tension); but in the enlarged heart, while each individual fiber exhibits the low active tension associated with heart failure, the increased mass of the enlarged heart and the decreased compliance of pathological origin allows the bearing of increased tension in the walls (via the law of Laplace). However, the analogy is pitched to mere description, and leaves the causal mechanisms in both the balloon and heart situations unaddressed. Medical students then fill the causal vacuum by employing the convenient assumption that the stretching in the heart causes the failure, as it causes the failure in the balloon. As was indicated earlier, activational problems are responsible for heart failure; enlargement (mistakenly likened to overstretch) is merely a later consequence of a variety of processes which are associated with failure. Correlation can easily be mistaken for causation in this case because causation is not represented in the analogy.

6. Magnification to the wrong grain size. An important aspect of the topic domain is missed because the analogy is pitched at the wrong level of magnification or

elaboration. It is not that the source domain does not address the aspect, as in #2, but rather that the analogy is cast in such a way that the relevant aspect is not noticed.

EXAMPLE: Certain applications of the Law of Laplace to lung function are demonstrated using an analogy to soap bubbles. In particular, soap bubbles are used to demonstrate what does not happen in the lungs (due to an agent secreted by the lungs called surfactant): Little air sacs of the lung do not empty into larger air sacs, as happens when small and large soap bubbles are in communication. However, it is difficult to understand why this does not happen in the lungs without knowing why it does happen in soap bubbles. And it is here that grain size is a problem. You would need to magnify considerably the structure and dynamics in the walls of the soap bubble to appreciate what happens differently in the lungs (due to the surface tension lowering effects of surfactant). The magnification would focus on the nature of the wall as two concentric circle layers with fluid in between them. In other words, the important lesson is in the "fine grain" of the soap bubble wall rather than in the typical, grosser image of soap bubbles connected by a pathway.

7. Misleading properties derived from common language meanings of technical terms. Ordinary language concepts are often employed analogically as technical terms. Their everyday, "public" connotation is overextended to their technical use in the topic domain.

EXAMPLE: In its technical usage the "compliance" (stretchiness) of vessels in the cardiovascular system is important as it contributes to the vessels' ability to exchange potential and kinetic energy, to ensure smooth movement of blood under the influence of the pulsatile beating of the heart. Students develop misconceptions that reflect the adoption of analogies for compliance that are consistent with common, ordinary language uses of "compliance" and that clash with its technical meaning. These include analogies of compliance as a "giving-way" (surrender) of vessels under the onslaught of blood, and even analogies that reflect the vessels' willingness to respond to orders from the nervous system (to, for example, dilate or contract, Feltovich et al., 1986).

(This misconception type, as well as #8, might be considered to be based more in metaphor than analogy. However, because the metaphorical descriptions lead to the adoption of specific source exemplars the features of which are mapped to the topic, the designation of these misconception types as analogical seems appropriate.)

8. Misleading properties due to connotations of nontechnical descriptive language. Inappropriate analogies are induced by the loose use of poorly chosen, connotatively loaded nontechnical descriptors.

EXAMPLE: The arterial vessels and other sections of the circulation are active contributors to blood flow. During the period of the heart's ejection (systole), blood from the heart is stored under pressure in the distended walls of the arterial vessels, so that when the heart is no longer actively ejecting (diastole) the recoil of these elastic vessels continues to propel blood. Perhaps to make a contrast to the period of active ejection of blood from the heart, diastole is often referred to loosely in instruction as a period of "run-off" or "drainage." These terms do not convey active motive force but, rather, passive flow from a region of high to low

static energy density. Students tend to think of the arterial vessels as passive receptacles and also tend not to consider adequately the active propulsion of blood during diastole, viewpoints more consistent with a notion of passive "seepage" than with the reality of blood flow during diastole.

It is interesting to note that a similar misconception occurs on the venous side of blood circulation due to the use in instruction of such language as "pooling" of blood in the veins, a term that carries the connotation of stagnation, or lack of interaction with the surround. Students develop the notion that some major components of the circulation can (at least temporarily) be removed from the active stream of circulation. Components of the venous circulation take on characteristics of reservoirs, "pools" that can be augmented and drained according to circumstantial needs for blood delivery, but that are not continuous parts of the overall circulation of blood.

9. Multiply based misconceptions. A single analogy may in fact lead to a variety of misconceptions, each involving a different influence from the list above.

EXAMPLE: The analogy of a rowing crew is commonly used to represent the functioning of the contractile units (sarcomeres) of muscle fibers. This analogy captures some aspects of the topic domain of muscle function well, particularly anatomic aspects of force production by oar-like ratcheting elements within the fiber. At the same time, the analogy (a) Indirectly Misleads (#1) with regard to some aspects of muscle ultrastructural function (e.g., conveying erroneously that the force producing units act in synchrony), (b) Misses (#2) some aspects (e.g., sarcomere characteristics related to their width as opposed to their length), and (c) Exports (#3) yet other misleading aspects (e.g., the idea that the force-generating units can get entangled, and thus fail to produce force). Much more will be said about the strengths and weaknesses of this particular analogy in the later section on multiple analogies.

Summary of analogy-induced misconception. Several ways that analogy can induce misconceptions have been presented. Although they differ in many respects, they have in common the characteristic that students' understanding of a topic domain is too exclusively determined by properties of an analogous source domain. In the various ways that we have illustrated, the mental representation of the topic domain is reduced to the source domain. It is in that sense that the misconceptions involve Reductive Analogies. In the next section we address the influences that contribute to the ready acceptance of analogical reductions.

3.2.3.2 The Analogical Bias: Patterns of Psychological Support for the Uncritical Acceptance of Deficient Analogies

Where do maladaptively reductive analogies originate? We have observed three sources. In the first, teachers and textbook authors recognize the pedagogical value of analogies in introducing difficult concepts and employ a conventional analogy for just that introductory purpose (while providing disclaimers about the limitations of the analogy and descriptions of appropriate corrections). They then find that their students have simplified their understanding of later, more advanced treatments of the concept to just what was covered by the introductory analogy. The second source is teachers and textbook authors who themselves have misconceptions. These may be associated with mistaking conventional approaches to introducing

a complex concept for complete accounts of the concept. In the third scenario, learners independently adopt analogies that, despite being both unconventional and incorrect, are nevertheless seductive in some way.

However, regardless of their origin, reductive analogies all share certain features. All of the scenarios involve a bias to over-rely on analogies: in the first case by students/learners in their implicit strategies for assimilating externally provided instructional information, in the second case by teachers, and in the third case by students/learners as regards their self-generated cognitive strategies. A further feature of the reductive analogies that we have observed is that they seem to lull the learner into an unquestioning acceptance that leads to a durable entrenchment of the misconception. Why, then, are maladaptively reductive analogies so readily adopted and so durably held? We will offer a partial account of the factors contributing to this bias, including its motivational origins and the methods implicitly employed for its self-justification. We will take an example of the third kind, an inappropriate analogy that is frequently adopted spontaneously. Understanding of the sources of such especially maladaptive analogies will also provide clues to why reductive analogies seem to be held so uncritically; that is, why they are so seductive of the belief that a full and accurate understanding has been achieved.

Support for false analogy in the Starling/Length-Tension example. The length-tension relationship of a single skeletal muscle fiber (studied in the laboratory, isolated from other muscles it normally interacts with in vivo) is essentially irrelevant to Starling's relationship between ventricular volume and cardiac output for the whole heart in vivo. (See the earlier discussion of #3 in the typology of ways analogies induce misconceptions.) Thus, when people adopt the analogy, they are doing it largely on their own. Despite its irrelevance, students give no sign that they consider the analogy at all controversial. What contributes to this unquestioning acceptance? More specifically, why would students want the analogy to be a useful one, and what causes them not to be suspicious of it? The following is a partial set of contributing forces.

1. **Bolstering due to similar appearing objects.** Volume is like length; force that produces ventricular (cardiac) output is like tension.

2. **Bolstering due to similar appearing relations.** The graph of the relationship between cardiac volume and output (Starling) has a similar shape to that of the relationship between skeletal muscle length and tension: They both have ascending slopes to a plateau.

3. **Reciprocating effects of separate bolstering elements.** The extent to which the similar appearing objects in the source and the topic bolster the analogy (#1) is increased by the fact that those similar objects also have a similar pattern of relationship in the source and the topic (#2). In a circular fashion, each is used to increase the credibility of the other.

4. **Bolstering due to assumptions of ontological similarity.** Why does it sound so initially plausible that the collective function of muscle fibers in the whole heart would be analogous to that of individual, isolated muscle fibers? Here we have what can fairly

be described as a fundamental ontological assumption about the natural world, namely that the world is structured in such a way that it is legitimately reducible by analogy: Wholes are like parts. The relational functioning of collections of fibers will be analogous to that of individual fibers (even though the objects involved, hearts and heart muscle cells, are markedly dissimilar): They will have the same shaped curve mapping their size aspects (i.e., length/volume) to their force aspects (i.e., tension/output), and the latter relationship will have the same theoretical explanation invoking the same mechanism (in particular, the sliding filament theory of muscle contraction--see Coulson et al., 1986).

Note further that this ontological bias can be reached through presuppositions deriving from two different ways of theorizing about the world: organicism and mechanism (Pepper, 1942). Organicistic accounts are noted for their attention to the symmetry of relationships found in microcosms and macrocosms (here, wholes are like parts). Mechanistic accounts would reach the same conclusion via such presuppositions as additivity and atomism: The collective Starling function "sums over" individual fiber length-tension functions (additivity); and individual fibers function the same when placed in context (in vivo) as when they are studied in isolation (atomism; i.e., the primary units of analysis are the individual components of a system and not the whole system itself--a fundamental bias of mechanistic approaches).

5. Indirect bolstering by convenience of explanation. In students' learning about heart failure, the well known "sliding filament" account of the length-tension relationship has a plausible sounding (but incorrect) extension to volume-output relationships in the heart. In other words, you do not need to learn a new underlying mechanism. Furthermore, you get to use a well-learned concept; the situation is akin to that of the person trying to improve his or her vocabulary who learns a new word and then orchestrates opportunities to use it. The conveniently overextended sliding filament theory then invokes a theme of overstretching for heart failure which (misleadingly) corresponds to a prominent and well known clinical sign of heart failure, enlargement.

6. Indirect bolstering by elimination of potential analogy blockers. We have seen that activational/energizing operations are neglected, in favor of mechanical overstretching operations (sliding filament theory), in the account of heart failure commonly preferred by medical students and by some physicians and popular textbooks. This elimination of activation cleans up a bit of potential untidiness in adopting the "length-tension: volume-output (Starling)" analogy. The absence of an activational component reinforces the idea that only mechanical factors are involved, and are involved in the same way, in the production of the two curves. That is, if heart failure is a mechanical phenomenon, and it is caused by "falling off the Starling curve," then mechanical rather than activational factors account for Starling. This consideration of Starling as being entirely mechanical makes it easy to ignore fundamental differences between the length-tension relationship and the Starling relationship that are due to the role of activation. Without an understanding of the differential roles of activation in the two curves, the "length-tension: volume-output" analogy gains strength.

But how does activation become neglected in the first place? One contributing factor is presented in #7.

7. Bolstering due to shared name. Degree of activation of muscle is not an issue if cardiac muscle is taken to be like skeletal, since skeletal muscle fibers always function at full activation (tetanic activation). So, if cardiac is like skeletal, variable activation is not an issue. Unfortunately, in this respect cardiac muscle fibers are not analogous to skeletal muscle fibers--the former are in fact variably activated (twitch activation).

So, the analogy between length-tension and volume-output is supported by the overreliance on the analogy of skeletal and cardiac muscle, which in turn is supported by the implicit (and reductive) bias that is described in the header of this sub-section and may be roughly instantiated in the present case as the belief that "all muscles are alike."

This neglect of activation in the cardiac volume-output relationship of course contributes to the mistaken acceptance of the "length-tension: volume-output (Starling)" analogy: without activation as an alternative account of volume-output and heart failure, volume-output is left with no competition for the mechanical overstretching account derived from the length-tension relationship.

8. Bolstering due to overall convergence of support. When many factors converge to support an analogy, as in the present "length-tension: volume-output (Starling)" example, the analogy is bolstered to a greater degree than if it had fewer supporting factors. "If so many things point to it, it must be right."

3.2.4 Summary

A theme of our research has been the recurring empirical observation of a pervasive tendency in cognition and instruction towards oversimplification of complex concepts (Fellovich et al., 1987; Spiro, et al., 1988; Spiro, et al., 1987). This tendency, which we have referred to as the "Reductive Bias," has been identified in both biomedical and historical domains and has been shown to take a great variety of forms. In this chapter we have confined our inquiry to just one cognitive arena, analogical learning, where we have provided several illustrations of the reductive force of analogies and the ramifications of that reductive force in the development of fallacious knowledge. We have shown in this chapter that (a) single analogies induce misconceptions involving the reduction of the topic domain to the source, and (b) there are certain mutually reinforcing biases that increase the likelihood that such reductive analogies will be uncritically adopted. Eight different ways that analogy can induce misconception were identified, and each was illustrated by an important misconception commonly held by medical students who have already had courses in the relevant content area. If one adopts reductive analogies, they cause trouble in later learning; and one is likely to adopt them.

These findings clearly indicates that analogies must be used with great caution. Even when they are used judiciously to initiate learners into a difficult subject area with appropriate caveats about their limitations, reduction of the topic domain to the source domain appears to be a too common occurrence. There are serious hazards involved in the use of analogies. The employment of a single analogy for a complex concept may impede the acquisition of more advanced understandings of that concept and to engender misconceptions. It is hoped that our typology of ways that analogy can induce misconception can serve to alert learners and teachers

to these potential hazards. However, the reductive force of analogies appears to be so great that even very detailed warnings are probably not sufficient by themselves. Approaches to analogy use more appropriate for complex concepts are needed, one of which is suggested in Chapter 4.2: integrated sets of *multiple analogies*.

3.2.5 Footnote to Chapter

The main contributors to this chapter were Rand Spiro, Paul Feltovich, Richard Coulson, and Daniel Anderson.

Section 4. TOWARDS MORE SUCCESSFUL ADVANCED-STAGE LEARNING: COGNITIVE FLEXIBILITY THEORY AND INSTRUCTIONAL HYPERTEXT

This section of the Report is concerned with repairing the kinds of problems of advanced knowledge acquisition described in the studies of Section 3. The second and third main phases of the research program, a theory of the development of cognitive flexibility in ill-structured domains and an instructional hypertext approach based on that theory, are recounted in Section 4.

Chapter 4.1 Knowledge Acquisition for Application: Cognitive Flexibility and Transfer in Complex Content Domains

4.1.1 Overview

A theoretical orientation to learning and instruction in ill-structured knowledge domains is presented. The theory is especially concerned with the application of knowledge in new situations (knowledge transfer), rather than the mere reproduction of knowledge in the way that it was originally learned. It is argued that knowledge transfer in complex and ill-structured domains is centrally dependent upon "cognitive flexibility." According to the theory, greater flexibility in the representation of domain knowledge will result from approaches that promote highly interconnected rather than neatly compartmentalized or hierarchicalized mental representations; that represent knowledge in terms of multiple, rather than single, prototypes and analogies; that increase the emphasis on learning from cases, while qualifying and restricting the scope of application of abstract principles; and that rely upon situation-dependent schema assembly rather than the retrieval of a rigid, prepackaged schema. A nonlinear system of learning and instruction that promotes these requisite features of cognitive flexibility is presented. In the system, cases or examples in a conceptual "landscape" are criss-crossed in a variety of directions, along multiple dimensions. Data is presented on the facilitation of transfer by learning with texts designed along the "landscape criss-crossing" model of case-centered text presentations, with re-pairings of cases on subsequent readings.

4.1.2 Introduction

A fundamental tenet of all recent theories of comprehension, problem solving, and decision making is that success in such cognitive arenas depends on the activation and appropriate application of relevant preexisting knowledge. Despite the substantial agreement on this general claim, we know very little about the organization of background knowledge and the method of its application to the understanding of new situations when, because of a combination of the breadth, complexity, and irregularity of a content domain, formulating knowledge in that domain to explicitly prescribe its full range of uses is impossible. We call knowledge domains of this type ill-structured and contrast them with more routinizable knowledge domains that we refer to as well-structured.¹ What does one do when relevant prior knowledge is not already organized to fit a situation (as will frequently be true in ill-structured domains, by definition) and so must be assembled from different knowledge sources in memory? This is a problem of knowledge transfer. We address a crucial issue in transfer: How should knowledge be acquired and organized to facilitate a wide range of future applications?

The principle contentions developed in this chapter are:

1. that theories in the cognitive sciences have produced a far better understanding of cognitive process in well-structured domains (WSDs) than in ill-structured domains (ISDs);
2. that theories appropriate for WSDs are in many ways inappropriate for ISDs--that, in fact, optimal conditions of learning and instruction in the two kinds of domains are opposite in several important respects;
3. that one of the most serious problems with treating ISDs as if they were WSDs is an inability to establish a basis for knowledge transfer, for the application of preexisting knowledge to new situations;
4. that transfer in ISDs is best promoted by knowledge representations that possess the following features: multiple interconnectedness between different aspects of domain knowledge, multidimensional or multiperspectival representation of examples/cases, and allowance for various forms of naturally occurring complexity and irregularity;
5. that in order for knowledge structures to possess the characteristics described in 4, emphasis must shift from the retrieval of a precompiled schema to the assembly of a situation-sensitive schema from knowledge fragments--the features described in 4 and 5 are characteristics of what we call cognitive flexibility;
6. that the best way to learn and instruct in order to attain the goal of cognitive flexibility in knowledge representation for future application is by a method of case-based presentations which treats a content domain as a landscape that is explored by "criss-crossing" it in many directions, by reexamining each case "site" in the varying contexts of different neighboring cases, and by using a variety of abstract dimensions for comparing cases.

Our primary claim is that in situations where complete comprehension or solution paths are neither inherent in the task or its description (i.e., are not implicitly or explicitly prescribed), nor anticipatable in the natural course of earlier learning and concomitant knowledge representation (either by explicit mention or by generative rule), then the key factors affecting the success with which prior knowledge is used to improve performance in a new situation will be the flexibility with which the relevant prior knowledge is represented in memory, and the mastery or control the individual has over those flexible representations (the ability to recombine elements of the representations, reorder the importance of elements in different contexts, and so on).

When knowledge cannot be routinized, mechanized, or automatized, it must be flexibly controlled. And control is not solely a process issue, independent of content. Executive control strategies require flexible knowledge structures to operate upon. In turn, control over flexible representations will give an individual more control over a content domain; rather than

monolithic prepackaged knowledge structures constraining an individual to apply knowledge in a fixed and limited manner, the individual controls the knowledge; that is, a great variety of nonpredetermined ways to adapt knowledge to the task and content elements involved in the new situation are available. This chapter presents a theory of learning and instruction, of knowledge representation and application, for the flexibility-based control that enables transfer. The goal of the program of research that we discuss is the validation of a set of basic principles and related instructional practices that will allow students to better apply the knowledge they acquire from formal schooling to new, real-world cases--knowledge that is built for use, not for imitative reproduction in artificial school or laboratory settings.

4.1.3 Schema-Theoretic Knowledge Representation and the Problem of Transfer

It is hard to imagine a more valued intellectual ability than that of independent thought--the ability to "think for yourself" when applying the necessarily limited knowledge acquired in formal schooling and training to the wide variety of situations in which that knowledge is relevant--the ability to use one's knowledge flexibly and to efficiently adapt it to varying contexts. Of what value is knowledge if its potential for application is much more limited than the range of uses to which it needs to be put? Very little, most would agree.

Nevertheless, a class of very influential contemporary approaches to knowledge representation, schema, frame, and script theories (e.g., Rumelhart & Ortony, 1977; Schank & Abelson, 1977) have had four interrelated shortcomings related to the problem of transfer. (As we see later, many of the following criticisms also apply to another class of approaches to knowledge representation: prototype and exemplar theories; e.g., Rosch & Mervis, 1975).

First, because these kinds of knowledge structures are frequently prepackaged (precompiled), they tend to be overly rigid. As a result, they provide very little opportunity for adaptation to diverse contexts of use.

Second, these modes of representation tend to isolate or compartmentalize aspects of knowledge that, in use, need to be interconnected. Again, the result is a limited potential for transfer.

Third, they have frequently treated complex subject matter as if it were simpler than it really was; complexities that transfer depends on have been artificially neatened.

Fourth, they have often implicitly assumed that knowledge domains possess more regularity or consistency across cases of application than they actually do.

These four characteristics of the recently predominant modes for representing knowledge, antithetical to the flexibility needed for wide application, are essentially simplifying assumptions: It is assumed that knowledge domains are simpler and more regular than they in fact are; these assumptions lead to representational approaches that are more rigid and compartmentalized than they ought to be. Historically, the assumptions of simplicity and regularity are strategically understandable and justifiable: Progress in new fields comes much more slowly without simplifying assumptions. And there has been considerable progress in cognitive science. However, when progress has been as slow as it has been on so important a

topic as transfer, then it is time to drop those simplifying assumptions that are causing progress to be impeded.²

The simplifying assumptions of cognitive science apply as well to dominant modes of education. Simplification of complex subject matter makes it easier for teachers to teach, for students to take notes and prepare for their tests, for test-givers to construct and grade tests, and for authors to write texts. The result is a massive "conspiracy of convenience." To take one example of the consequences of this tendency, Spiro, Feltovich, and Coulson (in preparation) have identified more than a dozen serious errors in the concepts held by a majority of medical students tested. Each of the errors was connected to a different kind of cognitive or educational oversimplification.

The overall effect of the simplifying features of knowledge representation systems and instructional strategies mentioned previously is a leveling tendency, a tendency towards monolithic approaches. Understanding is seen as proceeding in essentially the same way across instances of the same topic. Our view is different: The conditions for applying old knowledge are subject to considerable variability, and that variability in turn requires flexibility of response. Monolithic representations of knowledge will too often leave their holders facing situations for which their rigid "plaster-casts" simply do not fit. The result is the often heard complaint of students: "We weren't taught that." By this, they mean that they weren't taught exactly that. They lack the ability to use their knowledge in new ways, the ability to think for themselves. Our research focuses on fostering the kinds of flexible knowledge representations that would free students and trainees from the limitations of having to use the information they receive in instruction in only that way in which it was originally instructed (rigid knowledge structures that need to be applied rigidly).

To achieve this goal we systematically depart from the four interrelated simplifying assumptions of the schema-type theories that we identified earlier (which, again, are also common assumptions in schooling).

- Rigid, monolithic, prepackaged knowledge representations are replaced by flexible representations in which fragments of knowledge are moved about and assembled to fit the needs of a given context of application. Instead of prepackaged schemata, purpose-sensitive situational schemata are constructed, thus allowing knowledge to be used in different ways on different occasions for different purposes. The emphasis is shifted from prepackaged schemata to the ingredients for many potential schemata; schema selection is devalued in favor of schema assembly; storage of fixed knowledge is devalued in favor of the mobilization of potential knowledge. One cannot have a prepackaged schema for everything.

- We replace highly compartmentalized knowledge representations with structures characterized by a high degree of interconnectedness. Appropriate compartmentalization of knowledge in one situation may not be appropriate in another. Multiple interconnectedness permits (a) situation-specific categorization, (b) multiple access routes to relevant case precedents in long-term memory from the details of new cases, and (c) the development of a reservoir of potential analogies when case precedents are less literally relevant.

- Instead of inappropriate simplification, we work with complex subject matter (e.g., historical topics, military strategy, biomedical concepts), acknowledging and teaching towards the complexity inherent in them. (Later in this chapter we address the issue of how learners can be aided in managing difficult complexities.) And just as subject matter is frequently more complex than is realized, so, too, are the real cases to which knowledge about that subject matter has to be applied. Cases or examples must be studied as they really occur, in their natural contexts, not as stripped down "textbook examples" that conveniently illustrate some principle. The application of knowledge in new situations will be thwarted if the contingencies for application are more complex than the knowledge brought to bear.

- We relax the often unrealistic assumption of regularity, of routinizability, of well-structuredness--the assumption that subsets of individual cases (applications, examples) in a knowledge domain are sufficiently alike that they may be covered in common by a self-consistent system of general principles or rules. Thus, where schema-type theories tend to be overly general and to abstract away from individual cases that are classified in the same way, our approach sticks closer to the specific details and characteristics of individual cases. The consequences of treating ill-structured material as if it were well-structured are knowledge representations that are inapplicable in transfer situations or that produce inaccuracy. (Note that although complexity is highly related to irregularity, they are not the same thing. For example, the facade of the Old Executive Office Building next to the White House has a highly complex but regular structure. We work with domains that are complex and ill-structured.)

4.1.4 A Note on Related Research

Our approach has many connections to prominent thematic strands of current theoretical interest. However, the manner in which we instantiate these strands, especially the way they are configured in our overall approach, is novel. Therefore, because the main thing in common between what we do and the research of others involves overlap in subsets of the constituent themes emphasized rather than in the specifics of how those themes are applied and combined, we simply point to those thematic areas of overlap here: analogy (Carbonell, 1983; Gentner & Gentner, 1983; Gick & Holyoak, 1980, 1983; Rumelhart & Norman, 1981; Winston, 1983); learning from examples (Anderson, Kline, & Beasley, 1979; Carnine, 1980; Medin & Schaffer, 1978); reminding (Kolodner, 1980; Ross, 1984; Schank, 1982); connectionism, parallel distributed systems, "society of mind" (McClelland & Rumelhart, 1985; Minsky, 1979; Waltz & Pollack, 1985); tutorial guiding (Collins & Stevens, 1983); constructivity (J. Anderson, 1983; R. Anderson, 1977; Barsalou, in press; Bartlett, 1932; Spiro, 1977, 1980a, 1983); cognitive complexity (Scott, Osgood, & Peterson, 1979; Streufert & Streufert, 1978); automatic versus controlled processing (Shiffrin & Schneider, 1977; Schneider & Shiffrin, 1977); efficiency and cognitive economies in mental processes (Spiro, 1980b; Spiro & Esposito, 1981; Spiro, Esposito, & Vondruska, 1978); metacognition and learning to learn (Brown, Campione, & Day, 1981); perceptual and experiential aspects of conceptual memory (Bartlett, 1932; Spiro, 1983; Spiro, Crismore, & Turner, 1982); contextual encoding variability (Smith, Glenberg, & Bjork, 1978); linear separability (Medin & Schwanenflugel, 1981); fluid task environments (Lenat, Hayes-Roth, & Klahr, 1983); structural analysis of content-area text (Britton & Black, 1985; Spiro & Taylor, in press).

A good example of a researcher whose work incorporates many of the same themes that we do is Barsalou. As we are, he is centrally concerned with contextual variability in conceptual structure, cognitive flexibility, and temporary concept construction under contextual constraint (Barsalou, in press). Even here though, the similarities are not as instructive as they might at first appear because of the difference in the kind of domains studied. Barsalou works with relatively low-level concepts and categories. There are many fundamental differences in both internal structure and the cognitive processes that must operate on that structure between concepts like "Birds" and "Fruit" (or even "Places to Go On A Vacation"), on the one hand, and complex topics such as "20th-Century History" or "Military Battles," on the other hand. There are also big differences between semantic memory tasks and the application of knowledge to real cases. All the problems with concepts and categories addressed by researchers like Barsalou (e.g., graded structure; contextual variability; see also Medin & Smith, 1984; Smith & Medin, 1981) are so greatly exacerbated by dealing with knowledge application in real-world content domains that are made up of many concepts, that new issues of learning, representation, and transfer are inevitably introduced. The theoretical and empirical solutions corresponding to these issues produce a picture that overlaps very little between simple concepts and complex topical domains.

So, as promised earlier, because of the novel way we apply the preceding headlined themes, we move on to further discussion of the nature of our own approach. However, before proceeding we should explicitly acknowledge the most pervasive influence on our theoretical orientation and empirical procedures, the later work of Wittgenstein (1953). Our treatment of complex topical knowledge was inspired by prominent Wittgensteinian metaphors for knowledge organization and learning, especially the metaphor of the "criss-crossed landscape." These and other debts to Wittgenstein (e.g., the reliance on approximate processes of family resemblance and the role of visual-perceptual forms of representation in attaining synoptic conceptual understandings) will be obvious in the following section and throughout the chapter.

4.1.5. Ill-structuredness, Learning from Cases, and the Instructional Metaphor of the "Criss-Crossed Landscape": A Prescription for Transfer

4.1.5.1. Ill-structuredness

In many domains, if one compares the features of large numbers of cases, a subset of the following conclusions may be drawn:

- There are no rules or principles of sufficient generality to cover most of the cases, nor defining characteristics for determining the actions appropriate for a given case.
- Hierarchical relations of dominance and subsumption are inverted from case to case.
- Prototypes tend to often be misleading.
- The same features assume different patterns of significance when placed in different contexts.

- An explosion of higher order interactions among many relevant features introduces aspects of case novelty.

It is such failures of general principles as these that we designate by the family label ill-structuredness. As will be seen later, this failure of general principles in ill-structured domains is directly related to the most sweeping recommendation that will emerge from our theoretical orientation: an emphasis in learning and instruction upon multiplicity. Instead of using a single knowledge structure, prototype, analogy, and so on, multiple knowledge precedents will need to be applied to new situations (multiple schemas, several past cases, overlapping analogies). Under conditions of ill-structured complexity, single approaches provide insufficient coverage.

"Ill-structured Complexity in the Domain 20th-Century History." How would the events and phenomena of the 20th century be classified and presented in typical instruction? Some abstract system for organizing material would be developed. One likely approach would be to identify several themes of the 20th century to serve as organizing compartments (chapters, subheads), and then individual examples would be slotted where they seemed to best fit. "Moral Relativism" and "Knowledge Specialization" might both be discussed in a chapter on "Fragmentation as a Theme of 20th-Century Life"; a chapter about "Irrationality and Uncertainty in 20th-Century Ideas" might refer to "Freud's Psychological Theory of Unconscious Control of Motivation and Behavior" and "The Uncertainty Principle in Physics"; there might be a chapter on "Alienation" that would cite "Massive Bureaucracies Dwarfing the Individual" and "Mass Media Brainwashing"; and so on. Notice the implicit assumption of well-structuredness: Cases are taken as (often interchangeable) instances or illustrations of abstract themes. "The Advent of Nuclear Weapons" and "Existentialist Philosophy" both illustrate irrational aspects of the era, and it is the latter, more abstract, point that is taken as the important lesson.

However, think some more about the samples. Is not existentialist philosophy an example of alienation as well as a demonstration of the recent trend away from Cartesian views of rational man? Are there not advantages to grouping Freud's theories with the subliminal influences of mass media to bring out a point about how influences we are not aware of have come to increasingly control 20th-century life? In fact, any single organizational scheme for presenting examples from 20th-century history will suffer from two important shortcomings: (a) It will not be possible to present together all the cases that it would be instructive to present together, and cases that are not in the same physical vicinity in text will tend not to be closely connected in memory; and (b) much of the multifacetedness of the individual case will be lost as its significance is narrowed to the abstract point that the case is presented as an illustration of. Similar problems would arise if some other system were used to represent the 20th-century cases. For example, if events were presented in chronological order, connections between temporally distant events would be likely to be missed. Similar consequences would ensue from a division along lines like politics, economics, and culture. Any single system for organizing material would not establish enough connections along enough dimensions to prepare one to deal with the great varieties of discourse about the 20th century that a learner is likely to encounter.

4.1.5.2 The Importance of Cases

It is partly because of such ill-structuredness that classroom instruction in professional fields (e.g., medicine, business) is so often augmented by considerable case experience (and that instruction in basic, nonapplied fields of knowledge ought to be). Lectures and textbooks tend to stress generalizations, commonalities, and abstractions over cases. Such approaches are clearly very effective when cases tend to be pretty much alike in how they have to be analyzed and responded to (i.e., in well-structured domains). Unfortunately, it is too often true that such assumptions are unwarranted, and the greater convenience associated with the traditional pedagogies is negated if they are ineffective.

How is successful performance possible in the absence of generalizability across cases? The answer is by focusing analysis and knowledge representation more at the level of the individual case (example, occasion of use, event, or other cognate terms) and by guided experience with large numbers of cases. If cases come in many forms, one needs to see many cases in order to represent their varieties of contextual influences and configurations of features. In a well-structured domain, cases (examples, etc.) are luxuries, helpful in illustrating general principles and then discardable; the principles can stand for all their subsumed, interchangeable cases. In ill-structured domains, crucial information tends to be uniquely contained in individual cases--examples are not just nice, they are necessary. Our approach to knowledge acquisition is highly case-based.

The problem is compounded, however, by the fact that ill-structuredness has a limiting effect on case-based training as well as on the more abstract approaches: Just as there are no homogeneous systems of principles or rules that can be generally applied, there are no generally applicable prototype cases. Ill-structuredness means that there cannot be any recourse to homogeneity, to any single course of action across instances, whether it involves a single guiding principle, a single organizational scheme, or a single prototype case. (Of course, in relatively well-structured domains like "Trips to Fancy Restaurants," a single conceptual macrosystem can be very effective.) Real-world cases tend to possess a multifaceted complexity and thus need to be represented in lots of different ways in order to bring out those multiple facets. Then, instead of a single case being the basis for case-based cognitive processing, aspects of different cases need to be combined, and it is the resulting assemblages, made up of fragments of different cases, that underlie an important part of case-based reasoning. The reconstruction of knowledge requires that it first be deconstructed--flexibility in applying knowledge depends on both schemata (theories) and cases first being disassembled so that they may later be adaptively reassembled. (Of course, some integral, nondecomposed case information must also be retained for guidance in ecologically realistic case assembly.)

Thus, two kinds of flexibility are needed in knowledge structures for ill-structured domains, each of which is a central feature of the mental representations we attempt to build:

1. Each complex real-world case needs to be decomposed and represented along many partially overlapping dimensions (i.e., the same information must be represented in lots of different ways).

2. Many connections must be drawn across the decomposed aspectual fragments of the cases in 1, thus establishing many possible routes for future assembly and creating many potential analogies useful for understanding new cases or making new applications; it is for this reason that our instructional system emphasizes connections between apparently dissimilar cases (and aspects of cases)--connections among obviously similar cases are much more likely to be noticed without special training, and thus not cause transfer problems.

4.1.5.3 Nonlinear and Multidimensional Arrangement of Cases: The Metaphor of the Criss-Crossed Landscape

An example of how the two preceding theoretical commitments imply instructional methodologies is the use in our program of rearrangements of case presentation sequences so that the same case occurs in the context of various other cases. This enables different features of the individual multifaceted cases to be highlighted, depending on the characteristics of the other cases they are juxtaposed to, while simultaneously establishing multiple connections, including distant ones.

The discussion of case rearrangement brings us to the dominant metaphor employed in our theory of flexible knowledge representation and case-based instruction. Following the lead of Wittgenstein (see, for example, the preface to the *Philosophical Investigations*), we think of an ill-structured knowledge domain as akin to a landscape. Landscapes are often complex and ill-structured. No two sites are exactly alike, yet all sites possess many (but not all) of the salient features of the total landscape. The same could be said of landscape regions made up of several sites. The best way to come to understand a given landscape is to explore it from many directions, to traverse it first this way and then that (preferably with a guide to highlight significant features). Our instructional system for presenting a complexly ill-structured "topical landscape" is analogous to physical landscape exploration, with different routes of traversing study-sites (cases) that are each analyzed from a number of thematic perspectives.

The notion of "criss-crossing" from case to case in many directions, with many thematic dimensions serving as routes of traversal, is central to our theory. The treatment of an irregular and complex topic cannot be forced in any single direction without curtailing the potential for transfer. If the topic can be applied in many different ways, none of which follow in rule-bound manner from the others, then limiting oneself in acquisition to, say, a single point of view or a single system of classification, will produce a relatively closed system, instead of one that is open to context-dependent variability. By criss-crossing the complex topical landscape, the twin goals of highlighting multifacetedness and establishing multiple connections are attained. Also, awareness of variability and irregularity is heightened, alternative routes of traversal of the topic's complexities are illustrated, multiple entry routes for later information retrieval are established, and the general skill of working around that particular landscape (domain-dependent processing skill) is developed. Information that will need to be used in a lot of different ways needs to be taught in lots of different ways. Real cases (events, uses, etc.) have multiple slants, and, because the goal of widespread transfer must be to be ready for anything that realistically is likely to arise, learning and instruction must anticipate using many of these many slants in the ways they tend to occur.³ Criss-crossing a topic in many directions serves this purpose. It builds flexible knowledge. Accordingly, we construct nonlinear and multidimensional acquisition texts.

Illustration of findings in tests of the "Criss-Crossed Landscape" Hypothesis. One experiment that tested the hypothesis that case-centered text presentations, with re-pairings of cases on subsequent presentations, would promote novel transfer is described in this section. The learning materials were texts describing representative and influential cases of twentieth century phenomena (e.g., Freud's theory of unconscious motivation, uncertainty principle in physics, mass production and consumption). The control group read the text with a conventional hierarchical text organization: Each case was nested under one of the three abstract thematic chapter headings, such as "Irrationality and Disorder in the twentieth century." This seductive nesting of examples under an abstract point or general principle was intended to parallel the unidimensional treatment of cases and the abstraction-centered organization typical of textbooks and training manuals. The control condition thus simplifies, it presents the same case material according to a single well-structured system ("textbookization"), and it minimizes case interconnectedness across the compartments of classification employed in the acquisition text. In the experimental condition, the abstract "chapter" headings were eliminated; instead, the relationship among the cases was emphasized. To accomplish this, commentaries that linked the pairs of cases were written. (This extra material in the commentaries of the flexible groups was included in the section introductions and summaries of the control condition, to insure comparability of the content that both groups read.) High school junior and seniors, randomly assigned to conditions, served as the subjects. The control condition read the text organized the same way on all three presentations (akin to studying a course text for a test), but the experimental condition read re-paired cases and new commentaries linking the cases on each of the days.

There were two different types of tests. The first type of test measured *reproductive memory* for the information contained in the text, the second measured transfer. On the memory test, the control condition improved much more compared to their pretest performance than did the experimental condition (83.4% vs. 41.9%). This is not surprising: the control condition is expected to have mnemonic advantages that result in the rapid and accurate memorizability of the material presented in the acquisition phase of the experiments (compared to experimental conditions)--clear-cut "scaffoldings" are most effective for supporting reproductive, fact-retention-type memory (Anderson, Spiro, & Anderson, 1978; Ausubel, 1968). However, it is also expected that control conditions will produce overly inflexible, closed down representations that will result in less transfer than the experimental conditions. Such was the case in this experiment. The amount of improvement for the experimental condition on a transfer measure (gist memory for new, unstudied quotes of writers about the twentieth century -- the ability to use the instructed knowledge to enhance performance on related but uninstructed material) during the course of the experiment was *much* greater than for the control condition, 129% gain for the experimental condition compared to only 28% for the control condition!

The same interaction trend (control condition superiority in tests of memory for presented material and experimental condition superiority in ability to use the material in new situations) was observed in several other similar experiments, although the magnitude of the differences is often small. In all short experiments of this kind, the superiority of experimental conditions is less than it would be expected to be in longer studies (paralleling the time parameters in actual curricula), because the highly interconnected representations

produced by the experimental manipulations require some "critical mass" of information before the assembly processes they depend on can begin to operate at full effectiveness. With more longitudinal studies, there would be longer study time to process the complex materials and, especially, to permit more re-pairings of the cases, more criss-crosses of the topical landscape.

Results such as these suggest that there is a fundamental choice in methods of learning and instruction. Conventional methods seem to produce superiority when measured by conventional tests that stress reproductive, fact-retention types of memory. The methods developed from Cognitive Flexibility Theory are not as successful at producing mindless, imitative recall. However, if one agrees that the goal of learning and instruction should be the acquisition of generative knowledge with wide application in novel but partially related contexts, then it would seem that methods like ours are far preferable to the conventional ones.

4.1.6 Concluding Remarks

We have discussed various issues in an area of cognition and instruction that are currently poorly understood: how to get people to independently go beyond their specifically instructed knowledge. In response to this gap in knowledge, we offered a theory of case-based learning for transfer in ill-structured knowledge domains and suggested related methods of case-based instruction to produce flexible knowledge representations. Empirical indications of the success of the approach in producing transfer were presented.

From a practical point of view, the approaches we propose are not always easy, and they may result in some increases in the time and effort required in initial instruction in a domain. However, we also expect that that investment will be more than justified by the fact that it will not be wasted instruction, limited in application to situations that happen to fit some narrow range of explicitly established preparedness. Wasted either because the instruction was too narrow in the case prototypes used, thus limiting applicability to resemblant cases; or wasted because oversimplified general principles and rules were taught, accounting for too little of the relevant variability in the knowledge domain. Systems of instruction must be developed that produce knowledge that can be flexibly adapted to the wide variety of new situations to which it will need to be applied, even at some additional early cost.

We know of no area of human endeavor that lacks an ill-structured aspect. Success in ill-structured areas tends to come only with a considerable accumulation of actual case experience. Application of the learning principles we have proposed has the potential to take material that is either taught poorly or not taught at all (and thus left to the vagaries of haphazard acquisition from "experience" over long periods of time) and, for the first time, make that material directly instructable.

4.1.7 Footnotes for Chapter

The principal contributor to the early parts of this chapter was Rand Spiro; other contributions were made by Walter Vispoel, John Schmitz, Ala Samarapungavan, and Angela Boerger. The principal contributors to the last sections of this chapter were Aydin Durgunoglu, Rand Spiro, Angela Boerger, and John Schmitz.

¹ These terms receive further explication later in the chapter and are clarified in a more general manner by their use throughout the chapter--as "family resemblance" concepts, part of their definition can be no more than implicit in a complicated network of similarities and differences across uses (Wittgenstein, 1953).

² It should be noted that the relatively harsh tone taken toward schema-type theories throughout this chapter should be understood in the qualified sense in which it is intended. Rather than a blanket condemnation of these theories, we intend only to point out their shortcomings in enabling transfer in certain fairly common situations characterized by irregular complexity. Much of our earlier work has been in the schema theory tradition, and we are well aware of the importance of schema-type approaches. Again, our claims against those approaches are limited, namely to aspects of knowledge domains that are "ill-structured" (keeping in mind that even when well-structured knowledge is involved, the context of its application is frequently ill-structured). The next step after the kind of work that we describe in this chapter will be to combine schema-type representations with the more flexible kind we are developing, because the two approaches seem to be natural complements, possessing compensating strengths and weaknesses. Also, the fact that we emphasize the less studied ill-structured domains does not imply that we discount the existence of domains with substantial regularity--they are just not the focus of our research.

³ It is not necessary that criss-crosses of the landscape be exhaustive to produce generative structures. Rather, if enough of the topical landscape is portrayed, accurate anticipations of the structure of nonportrayed aspects will be possible. This can be seen by analogy to a situation in which one is presented with pictures of a man's face only when he has expressions of anger or happiness; one would be able to fairly accurately anticipate what that person would look like if he were sad, grateful, sleepy, and so on.

Chapter 4.2 An Example of the Cognitive Flexibility Theory
Principle of Multiple Representations:
Multiple Analogies as Antidotes for Analogy-Driven Misconception

4.2.1 Overview of Chapter

In order to combat the tendency toward oversimplification-based learning error that we have found to be associated with the use of a single analogy, an alternative approach involving integrated sets of multiple analogies is offered. Multiple analogies to correctly, yet manageably, capture the complexity of difficult concepts. In the multiple analogy approach, additional analogies are introduced that correctly convey information that is incorrectly represented (or not represented at all) in an earlier analogy. Thus, the pedagogical strengths of analogies are retained while their weaknesses are mitigated. The multiple analogy approach is illustrated by the example of force production by muscle fibers. The additional cognitive load introduced by multiple analogy sets may be mitigated by the employment of composite images, with situation-dependent selective instantiation of productive aspects of the composite. Finally, we present a more detailed picture of the variety of ways that adding analogies can affect the earlier learning outcomes derived from previously encountered analogies. We develop a nine-part taxonomy of the functions of new analogies and modifications of old analogies in promoting understanding.

4.2.2 Introduction

In Chapter 3.2 we examined in detail the hazards of using a single analogy to teach a complex new concept. What should be done to counter the tendencies toward oversimplification that come with the use of analogy and that we have found lead to misconceptions? The reductive force of analogies appears to be so great that it is not enough merely to tell people what the limitations of an analogy are. When teachers or texts provide such caveats for an instructional analogy, the result over time tends to be the same: The analogical core is what is retained. The remedy that we propose is to combat the power of a limited analogy with another powerful analogy that counteracts the limitations of the earlier one. In general, the antidote to any kind of Reductive Analogy might be the use of appropriately integrated multiple analogies, designed to vitiate the effects of three generic types of shortcomings in the use of single analogies (for a more detailed treatment of such shortcomings, see the earlier section on Analogy-Induced Misconceptions): (a) information that is missing from the source, (b) information in the source that is misleading about the topic, and (c) information that is inappropriately focused in the source. In this section we discuss an approach for selecting, integrating, and managing multiple analogies to promote full understanding of complex concepts in advanced stages of knowledge acquisition.

The discussion of our multiple analogy approach uses as an example the concept of muscle function introduced earlier. When the topic of force generation in a muscle is taught using the analogy to rowing teams, we saw that the resulting knowledge contained the seeds for various misconceptions. Some could be attributed to what the analogy misses (e.g., reduction in width with elongation of a muscle fiber) and some to aspects of the analogy that mislead (e.g., the synchronicity of movement across rowers).

Our approach to using multiple analogies is simply to introduce new analogies targeted at emending the missing, misleading, or misfocused information contained in earlier analogies. Then, as we will see in the next subsection, the various analogies are integrated in a composite image that has selective instantiations of the correct and useful information found in each analogy, but suppresses the inappropriate information. In other words, our procedure permits the retention of each analogy's strengths while discarding its weaknesses.

4.2.3 A Multiple Analogy Set for Muscle Fiber Function

We begin with a brief treatment of some of the more essential features of muscle fiber function, with reference to the muscle fiber portrayed in Figure 18.

The contraction of a muscle ultimately is accomplished by micro-structures within the muscle called cross-bridges. Cross-bridges are tiny "arms" that reach from one part of the muscle (myosin) out to an adjoining muscle structure (actin). The cross-bridges deliver the pulling force that accomplishes the contraction of a muscle. During a shortening contraction (non-isometric) each cross-bridge head releases from its site on the actin and moves to the next slot, possibly just vacated by its neighbor, while another cross-bridge pulling elsewhere prevents backsliding. The actin structures are thus pulled in towards the center of the myosin filament (see Figure 19). Each release and rebinding is called a "cross-bridge cycle." After a cross-bridge releases and rebinds, the cycle repeats until muscle deactivation occurs. That is, cross-bridge cycling continues over and over unless activating material (e.g., calcium) and/or usable energy (ATP) needed to energize the releases and rebindings become limited. The duration of the contraction is dependent on the degree of availability of these materials. The strength depends upon the number of cross-bridges pulling on actin filaments at any instant during the contraction. The somewhat simplified description presented here does not do justice to the real complexity of the process. The mechanisms by which the repeated cross-bridge cycles are accomplished involve intricate electrostatic, chemical, and physical processes and exchanges.

The processes of muscle function are clearly quite complex. As we have already seen, much of this complexity is not captured by the rowing team analogy (although the analogy is very helpful for conveying some of the anatomical aspects). Some strengths and weaknesses of the rowing team analogy are listed in the left and right columns, respectively, of Figure 20. Analogies relevant to other aspects of muscle fiber function are found in Figures 21-25, each with the same two-column structure of strengths and weaknesses.

The missing, misleading, and inappropriately focused information found in the right column of each of the analogy figures guides the selection of additional analogies used to promote fuller understanding; new analogies are chosen which convey the correct knowledge not conveyed




Figure 18. Schematic representation of the sarcomere. The sarcomere is the functional contractile unit of the muscle cell (fiber). The sarcomere is usually thought of as the unit contained between two of the z divisions in the (a) diagram. Actually a great many actin filaments are attached to each z disc and a great many myosin filaments are interspersed among them. A schematic cross-section through a muscle fiber is presented in the (b) diagram. It shows a hexagonal array in which each myosin filament is surrounded by actin filaments and each actin filament is surrounded by myosin filaments.





Figure 19. Schematic representation of a sarcomere during shortening. The illustrations (a)-(d) represent the sarcomere at various stages of a shortening contraction, between the rest length (a) and the maximum shortened length (d). The sarcomeres in a muscle cell (fiber) fill the cell from end to end. When they all shorten at once, the whole cell shortens.



Rowing Crew

Figure 20. Force production to create movement (contraction): the rowing crew analogy

CAPTURES

- (1) Anatomy of force producers: the little arms
- (2) Nature of the movement of the force producers: back and forth, hitting a resistance
- (3) Lots of individual force producers

MISSES OR MISLEADS

- (1) Conveys synchronicity: idea that all producers act in unison
- (2) Conveys notion that oars can get tangled (e.g. if boat too short)
- (Locus of L-T Bug)
- (3) Misses actual nature of gross movement (see Figure 4)
- (4) Misses things related to Width (see Figure 6)


by the old analogies. For example, one of the misconceptions engendered by the Rowing Team analogy concerns the nature of the gross movement of the main structures in relationship to each other. To produce contraction the "oars" must in fact pull something toward the center. Figure 21 modifies the analogy in Figure 20 to correct this incorrect aspect.

However, the analogy in Figure 21 still has the problem that it misses the crucial notion of attachment of the cross-bridges to some structure (the actin; see Figure 18) that gets pulled toward the middle--oars, in contrast, slip through water. This shortcoming suggests the analogy of the Turnbuckle (Figure 22). This analogy captures the correct idea of something being pulled toward the middle, without any change in the length of the structure doing the pulling (the myosin) or the structure being pulled (the actin). It also conveys the notion that there is a limit to the shortening process, which has also been missed by the earlier analogies. However, the Turnbuckle analogy has nothing corresponding to the cross-bridge anatomy depicted so well by the Rowing Team analogy; hence the new analogy is not a complete substitute for the old one. It is for such reasons that multiple analogies must often be maintained, rather than new ones simply supplanting old ones.

Another failing of the Turnbuckle analogy is that it provides no useful guidance on the limits to lengthening of the fiber (an aspect that is important, for example, in understanding heart failure). Here the analogy of the Chinese Finger Cuffs (Figure 23) provides a (partial) corrective: The double-wound spiral structure of the muscle covering (the collagen and elastin) produces a length limit (a kind of "wall"). Here again, some of the elements that the new analogy is weak in conveying are covered by earlier analogies (e.g., the anatomical information about individual force producers covered well by the Rowing Team analogy and the limits to shortness covered by the Turnbuckle).

We have now progressed considerably from the initial Rowing Team analogy, but essential information has still not been conveyed. Most prominently, there is still nothing about the recruitment of force producers (rowers and oars) to act in any given contraction cycle; not all force producers need act on any cycle and there is a process of selection from those available. (Note also how this analogy might serve to correct a notion that all rowers must be rowing at any one time, which is conveyed by the Rowing Team analogy.) This variable recruitment aspect of the topic is conveyed by the Slave Ship analogy in Figure 24.

More could be added about how muscle cells work. For example, the metabolic and energetic "life processes" of the muscle's force producers have not been addressed (Figure 25). (These processes are important in understanding muscle activation, which, in turn, is important in understanding cardiac failure.) Furthermore, analogies for them are not even



2 Rowing Crews Facing Each Other in the Same Boat

Figure 21. The nature of overall gross movement of parts in relation to each other (pulling toward the middle): the analogy of two crews facing each other in a long boat.


CAPTURES

- (1) Notion that action tends to pull something toward middle (the water)

MISSES

- (1) Attachment to some structure that gets pulled toward middle

(Problem is that oars slip through water)



Turnbuckle


Figure 22. The nature of overall gross movement: the turnbuckle analogy.

CAPTURES

- (1) Notion that action pulls something toward middle, with no change in length of puller
- (2) Notion that there are limits to shortening

MISSES

- (1) Cross-bridges
- (2) Individual force producers
- (3) Weak on limits to lengthening
(see Figure 6)



Chinese Finger Cuffs

Figure 23. Structure and function of the casing (elastin) surrounding the force producers (i.e., sarcomeres are contained within the casing: the Chinese finger cuffs analogy).

CAPTURES

- (1) Anatomy of the covering:
the double wound spiral
- (2) Thinning with stretch (the
dimension missed in Figure 3

(Important in right-side
L-T function-squashing)
- (3) Limits to length
 - A. "Wall" beyond which
you can't go
 - B. At "wall," will break
before it lengthens

MISSES

- (1) Anything about force
producers, etc.
- (2) Anything about the stuff
inside the casing

(Where's the beef?)




Figure 24. Recruitment of force producers: the galley ship analogy.

CAPTURES

- (1) Control that selects and recruits which force producers are to work (on any stroke)
- (2) Recruitment aspects of activation

MISSES

- (1) Communication mechanisms
- (2) Anything about internal metabolic, energetic, life processes of the force producers
(The galley ship rowers)

easily created. (Note, however, that we have been talking about the force producers as people rowing. This may actually be a fairly good analogy for the processes under discussion here.)

At this point of the analysis a plateau of sorts has been reached. Many essential aspects of the topic have been covered. And if one considers that much of what medical students understand (and misunderstand) about muscle function is guided by the Rowing Team analogy, the advance in understanding easily enabled by the additional analogies is very substantial indeed. A complex topic that typically is oversimplified can now be more readily grasped.

4.2.4 Integrating Multiple Analogies

The most important step in the integration of multiple analogies has already been covered by the selection procedure discussed in the previous section. The multiple analogies in a set are interlocking, since each new analogy is chosen to correct the negative aspects (the right-columns in Figures 20-24) of the preceding analogies. The remaining discussion of integration is merely an unpacking of the integrative machinery implicit in the procedure for selecting analogies.

Once a set of analogies has been selected, such that the unproductive missing, misleading, or poorly focused information for each analogy has been modified or canceled out by some aspect of a succeeding analogy, then the multiple analogy set must be integrated. There are two main ways of adding new analogies to earlier analogies. These will have an impact on how the additions function to produce new understanding. (A more detailed taxonomy of the functions of new analogies and modification of old ones is presented in the next section. Although there are relationships between the functions of new analogies discussed there and the patterns of combination described here, they are distinct topics. We restrict ourselves here to the smallest number of types of combination necessary for introducing the composite imaging procedure in the next section.) In one of the two combination patterns, Modified Analogies, new analogies merely amend an earlier one (e.g., correcting the Rower analogy to have two sets of rowers pulling toward the middle; Figures 20 & 21). In the second way, New Analogies, a totally new analogy is added (e.g., introducing the Turnbuckle analogy following the Rower analogy).

The class of Modified Analogies contains two subtypes: Simple and Complex. With the addition of a Simple Modified Analogy there is no need to retain the old analogy; the new analogy amends incorrect information from the old one but loses none of its essential information. This type of Modified Analogy is the easiest to deal with. First, the modified analogy is merely substituted for the earlier one (e.g., replace the drawing in Figure 20 with that in Figure 21). Second, the critical new aspect of the substituted Modified Analogy is accentuated. For instance, the image of two sets of rowers facing each other and rowing toward the middle is marked for special notice. This may be accomplished, for example, by suggesting to students that they envision one of the sets of rowers as colored red and the opposed set as blue; or by placing eye-glasses on the rowing team members to accentuate the directions they are facing. The point, again, is to accentuate the features of the Modified Analogy that constitute the main change from the prior analogy.




Figure 25. The biochemical life processes of the force producers. How they get energy, nutrition, expend it, etc.--how they get "juiced up." This is the ATP energy cycle of the muscle and all its components.

In a Complex Modified Analogy, the analogy used for correction of the earlier analogy contains seeds of misconceptions that are best mitigated by maintaining an aspect of the earlier analogy. Here the prior analogy cannot simply be "updated." Analogies of this type do not occur in our muscle fiber example. However, integration would be handled by the same selective instantiation feature of our composite imaging technique that is illustrated with New Analogies, discussed next.

4.2.5 Making Multiple Analogy Representations Cognitively Tractable: Composite Images With Context-Dependent Selective Instantiation

How should one keep track of several analogies and the complicated pattern in which they are combined (i.e., the patterns by which the left-columns of Figures 20-24 are retained and the right-columns overridden as successive analogies are introduced)? The cognitive demands are considerable. To make the psychological management of this complexity more tractable, we introduce the technique of "Composite Imaging with Selective Contingent Instantiation" (CISCI). The remainder of this section is devoted to a description of the CISCI procedure. However, although conventional means of teaching concepts will not receive attention, it must be emphasized at the outset that the composite images should be perceptual adjunct aids that supplement rather than replace verbal instruction.

In the CISCI procedure, a complete set of images, each corresponding to one of the multiple analogies (other than those superceded by a Simple Modified Analogy), would be available in a latent state. That is, all of the images that are to be combined are represented as potentials in a state of readiness for instantiation. (We intend for these representations to refer both to latent mental images and to potentially be displayable external images, e.g., computer graphical displays. Thus our discussion refers both to mental representation and its parallel in instructional support systems.) Which of the latent images (analogies) of the total composite image are actually instantiated at any one time is aspect-dependent. That is, a component of the (latent) total composite image will be a part of an actual instantiated composite image as a function of the aspects of muscle fiber function under consideration. For example, in contexts where it is important to focus on the way muscle fiber is affected by longer lengths and length limitations, the Chinese Finger Cuffs image would become salient; when thinking is directed at muscle functioning under short lengths and toward the limits to shortening, the Turnbuckle image would be "programmed" for high salience. Component images that are not activated are suppressed; they may either not be visible (in "the mind's eye" or in an external support system) or, preferably, they may be instantiated very faintly, as a kind of background to the aspect-dependent focal images.

Finally, at the same time that either the Turnbuckle or the Chinese Finger Cuffs (continuing the last example) is relatively salient within the instantiated CISCI, images corresponding to the other analogies that are not being actively suppressed (i.e., that retain their usefulness for the conceptual aspect under consideration) would simultaneously be superimposed within the total image. So, if the Chinese Finger Cuffs are active, that image would be overlaid by the Double Rowing Team image; an image of two sets of rowers pulling toward the center would be superimposed on the cylinder of the finger cuffs, enabling the positive properties of both analogies to make a simultaneous contribution.

In sum, the characteristic of aspect-dependent selective and contingent instantiation of component images (along with the initial motivation and basis for selecting each of the analogies, that is, overcoming some inappropriate information in an earlier analogy--the right columns in Figures 20-24), makes it possible for understanding of the concept to be supported only by the correct and productive features of each analogy, with the incorrect and unproductive features cancelled out. The CISCi can be thought of as a montage of simultaneous ways that muscle fiber function can be "seen as" (Wittgenstein, 1953) something else. Parts of the montage are emphasized in visualization as a function of the productive left column information that needs to be stressed and the unproductive right column information that needs to be suppressed, given the conceptual aspects under consideration at any one time.

The CISCi method we have described is thus one of *nonabstractive integration*. By using a composite of several analogies, understanding would neither be abstractively reduced to a superordinate system nor to any one of the analogies acting as privileged with respect to the others. In contrast, in abstractive integration the individual elements are replaced by a subsumptive abstraction that stands for the elements and/or their combinations. By presenting the composite analogies as simultaneously overlapping images, the composite leads to a perceptual integration. Thus, as in the perception of complex objects like human faces, a whole is psychologically graspable without loss of information about parts--although one's image of the face of a close friend is typically perceived as a simultaneous physiognomy, the components of the face are recoverable (i.e., perceptual focus can shift to individual facial features, such as the size of the nose or the distance between the eyes). The physiognomy is perceptually integrated while the component parts continue to make their contributions individually (in contrast to the supplanting of individual elements in an abstractive integration).

The kind of perceptual composites or "integral visual displays" for representing nonreducible multidimensional information that we are proposing have been receiving increasing support in recent years (Chernoff, 1973; Spiro, 1982; Spiro et al., 1987; Spiro & Myers, 1984; Wainer & Thissen, 1981). The more complexly structured the domain, the more it becomes impossible in principle to conceptually chunk large groups of information; so, instead of a conceptual chunking of information to lessen the cognitive load, we recommend a perceptual chunking.

4.2.6 A Taxonomy of Functions of Additional Analogies in Promoting Conceptual Understanding

There are a variety of reasons that new analogies might reasonably be introduced to augment earlier analogies in promoting understanding. Earlier, for purposes of illustrating the mechanics of the composite imaging technique, we presented a very simple classification of the kinds of analogy additions that can be made. Here we offer a more detailed accounting of the ways that additional analogies can contribute to understanding. A preliminary and nonexhaustive classification of such functions of multiple analogy is given below (mostly carrying over examples from muscle fiber function).

1. **Supplementation (with New Analogy).** Aspects of a topic domain that are missed by earlier analogies are covered by a supplementary analogy. (Also see Burstein, 1983.)
EXAMPLE: Adding the Turnbuckle analogy to cover the information missing from the Rowing Crew analogy about attachment to some structure that gets pulled (Figure 22).

2. **Correction (with New Analogy).** Aspects of a source domain that mislead about the topic domain are corrected (without altering the correct information in the earlier analogy). **EXAMPLE:** Replacing the Tumbuckle with Chinese Finger Cuffs in the context of understanding muscle fiber function at long lengths (Figure 23).

3. **Alteration (of an Earlier Analogy).** Sometimes an incorrect element in an earlier analogy can be dealt with by modifying or "patching" it, rather than by fundamentally changing it or replacing it altogether. **EXAMPLE:** Modifying the Rowing Crew analogy to have the crew split in half and row towards the center (Figure 21).

4. **Enhancement (of an Earlier Analogy).** A refinement of an earlier analogy that deepens the understanding of the topic domain--that overcomes some earlier superficiality, without either altering the existing components of the earlier analogy (#3, Alteration) or changing its magnification (#5, Magnification). **EXAMPLE:** Having the cockswain call out the numbers of the rowers (recruitment of cross-bridges) to convey the notion of asynchronous movement and degrees of recruitment (Figure 24). Note that the Enhancement impresses an appropriate new layer of mechanism on the phenomenon. As a side effect, the Enhancement also repairs a misleading implication of an earlier analogy (Figure 20; first item in the right column), without repairing the earlier analogy itself. Thus it differs from #3, Alteration. In terms of the rules for constructing Composite Images, this is a Simple Modified Analogy, whereas Alteration is a Complex Modified Analogy.

5. **Magnification (or Elaboration).** An aspect of the topic domain that is not correctly captured because of the "grain size" of the analysis (and corresponding image) evoked by the source analogy is addressed by scale alterations (i.e., changes in grain size). **EXAMPLE:** While the introduction of the cockswain (Figure 24) captures the notions of recruitment of force producers and asynchronous movement, it is simply opaque regarding the biochemical life processes involved in variably energizing the force producers. Perhaps the best way to remedy this problem is by a many-times magnification of an individual rower in the rowing crew (i.e., "people," and their intake, breakdown, and utilization of energy, as an analogy for the metabolic aspects of the force producers).

Note that there are two possibilities following Magnification: introduction of a new analogy for the magnified section of the old analogy and retention of the old analogy on a different scale. An example of Magnification involving the retention of an earlier analogy but with change of scale was encountered earlier in the Law of Laplace example, where magnification of the soap bubble wall was a key step.

Of course, when a more synoptic view is required, grain size may be decreased in a "reverse magnification." Also, a function related to Magnification would be served by rotating the image of the source domain so that previously obscured elements are brought into view (e.g., as would happen if the first image involved the front of a thing and one needed to augment understanding by later looking at the back of the thing).

6. **Perspective Shift.** Complex domains can often be thought of in fundamentally different ways. In such cases, a new analogy may convey a different perspective than that

conveyed by earlier ones. **EXAMPLE:** The Rowing Crew represents well the perspective of movement production. The Finger Cuffs represent well the perspective of limits to movement.

A related use of new analogies for contributing to understanding involves the introduction of an analogy that addresses underlying causal mechanisms after an earlier analogy that was more concerned with the surface form of a situation. **EXAMPLE:** When heart failure is taken as analogous to an overstretched balloon that has lost its tensile strength, some aspects of superficial similarity are captured: two things have gotten big and have lost their ability to produce force. However, as we saw when this kind of reductive analogy was discussed earlier, despite the fact that the underlying causes are greatly different, students mistakenly assume they are the same. Introduction of an analogy that more correctly characterizes the cause of heart failure (e.g., a dying car battery) would make it less likely that students would mistake gross "input-output" correlation for causation in the development of the misconception about heart failure.

7. Competition. More than one analogy competes as an overall account of the same domain. One analogy eventually supercedes its competitor, which is discarded. Only one analogy occupies a specific analogical niche. **EXAMPLE:** A common analogy employed by medical students in understanding congestive heart failure is that of an overstretched, "bagged out" balloon incapable of generating force (Coulson et al., 1986; also see the earlier discussion of the false analogy between the length-tension relationship for individual muscle fibers and Starling's relationship between cardiac volume and output). Since heart failure is really due more to impotence within energetic activational factors of cardiac muscle than to anything related to mechanical overstretching, a more appropriate overall analogy would involve something like the dying battery (or perhaps a fouled spark plug or broken distributor wire) introduced in the last section.

8. Sequential Collocation. Successive stages in a process are each represented by an analogy. In other words, a single analogy is used for each identifiable "segment" of a phenomenon, and the analogies are simply collocated (without any integration being necessary).

The last two categories are the main ones addressed by Collins and Gentner (in press) in one of the few discussions of multiple analogies in the literature (see also Burstein, 1983). A brief discussion of their work on evaporation models will help to clarify our own approach. Again, Collins and Gentner deal with multiple analogies in two different ways. In the first, they divide up the water cycle into successive stages (e.g., what happens in bodies of water; what happens in the sky; and so on). This is #8, Sequential Collocation. A multiple analogy set would then be the group of analogies for each stage. Their second use of multiple analogies involves choosing the best analogy for each stage. That is, they discuss alternate analogies for each stage (e.g., alternate analogies for what happens in bodies of water), with the ultimate goal of selecting the most appropriate analogy for that stage. This is #7, Competition.

The fundamental difference from our approach should be clear: Our multiple analogies are all applied to the same stage of a process, and all of the analogies in the set are partially correct--they are all "best" at conveying something about the topic domain. More generally, we claim that for most domains of any complexity the "pieces" (e.g., stages, sectors, aspects, etc.) that the domain is divided into will each need to be covered by more than one analogy. And this

is what our "aspect-dependencies" do. However one compartmentalizes a complex domain, the pieces cannot usually be effectively treated with a single analogy. (A related issue that we have not addressed here concerns the extent to which "pieces" of a complex domain can be adequately treated separately from each other at all.)

4.2.7 Conclusion

Analogies have traditionally been seen as a powerful positive force in learning new material. We have shown a countervailing tendency: In more advanced forms of knowledge acquisition, the power of a single analogy is often negatively exercised, blocking advanced learning and leading to misconception. We have argued that the best way to counter this powerful negative force is with equally powerful positive forces--in order to combat the elements that induce misconceptions in an otherwise useful single analogy, import new analogies that powerfully convey the correct knowledge. Multiple additional analogies are chosen to correct the problems introduced by any single analogy, without cancelling the beneficial effects of the single analogy.

In order to lay the groundwork for such an approach, we have offered a detailed conceptualization of the diverse functions of multiple analogies in learning, a procedure for selecting and integrating multiple analogy sets so that they convey complexity without inducing misconceptions, and a composite imaging technique to lessen the cognitive load of working with a multiple analogy set. Taken together, these elements constitute a comprehensive program for the employment of multiple analogies in support of advanced forms of complex knowledge acquisition.

4.2.8 A Final Note: Simplicity and Complexity in Advanced Knowledge Acquisition

This chapter and chapter 3.2 dealt explicitly with the reductive force of single analogies in the development of misconceptions and with the role of multiple analogies in promoting correct understandings of complex concepts. However, the work on analogy has a more general purport as well. It can also be treated as a kind of apology about learning in complex domains. The detailed treatment of hazards of analogy use and remedies for them has close parallels in the use of any mode of cognitive support for complex new learnings. In the findings of our other studies, the maladaptive reductive force of single analogies is paralleled by misconception-inducing reductive forces of a single schema, single mode of organization, single line of argument, single precedent example, single prefigurative "world view," and so on. The antidote for these maladaptive forces of simplification is in each case the systematic assembly of *multiple knowledge sources*--integrated multiple analogies, compiled fragments from diverse schemata, re-presentations of the same information under different organizational schemes, multilinear lines of argument, and multiple precedent examples. So, it could be said that this work, besides being about analogy *per se*, is also, in an important sense, about the seductive force of all singular approaches to complex learning, the prevalent role of oversimplification in the development of entrenched fallacies of conceptual understanding, and the general importance of assembling multiple knowledge sources to support the mastery of complex concepts at advanced stages of knowledge acquisition.

4.2.9 Footnote to Chapter 4.2

The main contributors to this chapter were Rand Spiro, Paul Feltovich, Richard Coulson, and Daniel Anderson.

Chapter 4.3 An Introductory Rationale for Hypertext: Cognitive Flexibility, Constructivism, and Random Access Instruction for Advanced Knowledge Acquisition in Ill-Structured Domains

This chapter and the next three that follow all deal with a hypertext approach based on Cognitive Flexibility Theory, an approach that builds upon and incorporates the lessons learned from all of the research discussed earlier in this Report. The present chapter, provides background and a theoretical rationale that introduce the hypertext approach. The succeeding chapters present a detailed theory of hypertext design, illustrative examples from our hypertexts, and empirical findings of transfer success associated with the approach.

4.3.1 Introduction: The Complex Context of Learning and the Design of Instruction

A central argument of this chapter is that there is a common basis for the failure of many instructional systems. The claim is that these deficiencies in the outcomes of learning are strongly influenced by underlying biases and assumptions in the design of instruction that represent the instructional domain and its associated performance demands in an unrealistically simplified and well-structured manner. We offer a constructivist theory of learning and instruction that emphasizes the real-world complexity and ill-structuredness of many knowledge domains. Any effective approach to instruction must simultaneously consider several highly intertwined topics, such as

- the constructive nature of understanding;
- the complex and ill-structured features of many, if not most, knowledge domains;
- patterns of learning failure;
- a theory of learning that addresses known patterns of learning failure.

Based on a consideration of the interrelationships between these topics, we have developed a set of principled recommendations for the development of instructional hypertext systems to promote successful learning of difficult subject matter (see Spiro, Coulson, Feltovich, & Anderson, 1988; Spiro & Jehng, 1990). This systematic, theory-based approach avoids the ad hoc character of many recent hypertext-based instructional programs, which have too often been driven by intuition and the power of the technology.

In particular, we argue that

- Various forms of conceptual complexity and case-to-case irregularity in knowledge domains (referred to collectively as ill-structuredness) pose serious problems for traditional theories of learning and instruction.
- Cognitive and instructional neglect of problems related to content complexity and irregularity in patterns of knowledge use leads to learning failures that take common, predictable forms. These forms are characterized by conceptual oversimplification and the inability to apply knowledge to new cases (failures of transfer).
- The remedy for learning deficiencies related to domain complexity and irregularity requires the inculcation of learning processes that afford greater cognitive flexibility: this includes the ability to represent knowledge from different conceptual and case

perspectives and then, when the knowledge must later be used, the ability to construct from those different conceptual and case representations a knowledge ensemble tailored to the needs of the understanding or problem-solving situation at hand.

- For learners to develop cognitively flexible processing skills and to acquire contentive knowledge structures which can support flexible cognitive processing, flexible learning environments are required which permit the same items of knowledge to be presented and learned in a variety of different ways and for a variety of different purposes (commensurate with their complex and irregular nature).

- The computer is ideally suited, by virtue of the flexibility it can provide, for fostering cognitive flexibility. In particular, multidimensional and nonlinear hypertext systems, if appropriately designed to take into account all of the considerations discussed above, have the power to convey ill-structured aspects of knowledge domains and to promote features of cognitive flexibility in ways that traditional learning environments (textbooks, lectures, computer-based drill) could not (although such traditional media can be very successful in other contexts or for other purposes). We refer to the principled use of flexible features inherent in computers to produce nonlinear learning environments as Random Access Instruction (Spiro & Jehng, 1990).

Following our injunction to consider all crucial issues in the learning and instruction environment jointly, we will develop the following compound argument, which integrates the claims presented above:

- Characteristics of ill-structuredness found in most knowledge domains (especially when knowledge application is considered) lead to serious obstacles to the attainment of advanced learning goals (such as the mastery of conceptual complexity and the ability to independently use instructed knowledge in new situations that differ from the conditions of initial instruction). These obstacles can be overcome by shifting from a constructive orientation that emphasizes the retrieval from memory of intact preexisting knowledge to an alternative constructivist stance which stresses the flexible reassembly of preexisting knowledge to adaptively fit the needs of a new situation. Instruction based on this new constructivist orientation can promote the development of cognitive flexibility using theory-based hypertext systems that themselves possess characteristics of flexibility that mirror those desired for the learner.

In summary: Ill-structured aspects of knowledge pose problems for advanced knowledge acquisition that are remedied by the principles of Cognitive Flexibility Theory. This cognitive theory of learning is systematically applied to an instructional theory, Random Access Instruction, which in turn guides the design of nonlinear computer learning environments we refer to as Cognitive Flexibility Hypertexts.

Selective Focus on Advanced Knowledge Acquisition in Ill-Structured Domains

The argument developed in this chapter is not intended to cover all aspects of constructive mental processing. Similarly, instructional technology is a broad topic that will not be exhaustively addressed in this chapter. Rather, we will focus on a set of issues implicated by consideration of some special instructional objectives (Merrill, 1983) and the factors contributing to their attainment. In particular, we will be concerned only with learning

objectives important to advanced (post-introductory) knowledge acquisition: to attain an understanding of important elements of conceptual complexity, to be able to use acquired concepts for reasoning and inference, and to be able to flexibly apply conceptual knowledge to novel situations. Furthermore, we will consider only complex and ill-structured domains (to be defined below). This combination of ambitious learning goals and the unobliging nature of characteristics associated with certain knowledge domains will be seen to present special problems for learning and instruction that call for special responses at the level of cognitive theory and related instructional interventions.

We will argue that one kind of hypertext approach is particularly appropriate for this constellation of features associated with the instructional context. The omission of other varieties of computer-based instruction from our discussion does not imply any negative evaluation of their merits. Indeed, in other instructional contexts the kinds of hypertexts we will discuss would be inappropriate (e.g., computer-based drill would be better suited to the instructional objective of memorizing the multiplication tables; see Jacobson & Spiro, 1991b, for the presentation of a framework for analyzing instructional contexts to determine the choice of educational technologies).

In what follows, we wish to illustrate how a particular set of factors in the instructional context (including learning goals and the nature of the knowledge domain) and a set of observed learning deficiencies jointly lead to a recommended cognitive theory-based instructional approach.

4.3.2 The Nature of Ill-Structured Knowledge Domains and Patterns of Deficiency in Advanced Knowledge Acquisition

Ill-structured Knowledge Domains: Conceptual Complexity and Across-Case Irregularity

An ill-structured knowledge domain is one in which the following two properties hold: 1) each case or example of knowledge application typically involves the simultaneous interactive involvement of multiple, wide-application conceptual structures (multiple schemas, perspectives, organizational principles, and so on), each of which is individually complex (i.e., the domain involves concept- and case-complexity); and 2) the pattern of conceptual incidence and interaction varies substantially across cases nominally of the same type (i.e., the domain involves across-case irregularity). For example, understanding a clinical case of cardiovascular pathology will require appreciating a complex interaction among several central concepts of basic biomedical science; and that case is likely to involve differences in clinical features and conceptual involvements from other cases assigned the same name (e.g., other cases of "congestive heart failure"). Examples of ill-structured domains include medicine, history, and literary interpretation. However, it could be argued that even those knowledge domains that are, in the main, more well-structured, have aspects of ill-structuredness as well, especially at more advanced levels of study (e.g., mathematics). Furthermore, we would argue that all domains which involve the application of knowledge to unconstrained, naturally occurring situations (cases) are substantially ill-structured. For example, engineering employs basic physical science principles that are orderly and regular in the abstract and for textbook applications (Chi, Feltovich, & Glaser, 1981). However, the application of these more well-structured concepts from physics to "messy" real-world cases

is another matter. The nature of each engineering case (e.g., features of terrain, climate, available materials, cost, etc.) is so complex and differs so much from other cases that it is difficult to categorize it under any single principle, and any kind of case (e.g., building a bridge) is likely to involve different patterns of principles from instance to instance. Similarly, basic arithmetic is well-structured, while the process of applying arithmetic in solving "word problems" drawn from real situations is more ill-structured. For example, consider the myriad ways that arithmetic principles may be signaled for access by different problem situations and problem wordings.

Advanced Knowledge Acquisition: Mastery of Complexity and Preparation for Transfer

The objectives of learning tend to differ for introductory and more advanced learning. When first introducing a subject, teachers are often satisfied if students can demonstrate a superficial awareness of key concepts and facts, as indicated by memory tests that require the student only to reproduce what was taught in roughly the way that it was taught. Thus, in introductory learning, ill-structuredness is not a serious problem. Learners are not expected to master complexity or independently transfer their acquired knowledge to new situations. These latter two goals (mastery of complexity and transfer) become prominent only later, when students reach increasingly more advanced treatments of the same subject matter. It is then, when conceptual mastery and flexible knowledge application become paramount goals, that the complexity and across-case diversity characteristic of ill-structured domains becomes a serious problem for learning and instruction.

Patterns of Advanced Learning Deficiency in Ill-Structured Domains and Remedies in "Cognitive Flexibility Theory"

In this section we briefly review two related bodies of research: the nature of learning failures in advanced knowledge acquisition and new theoretical approaches to more successful advanced learning and instruction.

Forms of a "reductive bias" in deficient advanced knowledge acquisition. Advanced knowledge acquisition, that very lengthy stage between introductory treatments of subject matter and the attainment of expertise for the subject, has been very little studied (certainly in comparison to the large number of studies of novices and experts--e.g., Chase & Simon, 1973; Chi et al., 1981; Feltovich, Johnson, Moller, & Swanson, 1984). However, in our own recent investigations of advanced learning in ill-structured domains, we have found a number of notable results, some of which were somewhat surprising (Coulson, Feltovich, & Spiro, 1989; Feltovich, Spiro, & Coulson, 1989; Myers, Feltovich, Coulson, Adami, & Spiro, 1990; Spiro, Feltovich, Coulson, & Anderson, 1989). These results may be summarized as follows:

- Failure to attain the goals of advanced knowledge acquisition is common. For example, when students are tested on concepts that are consensually judged by teachers to be of central importance and that have been taught, conceptual misunderstanding is prevalent.
- A common thread running through the deficiencies in learning is oversimplification. We call this tendency the reductive bias, and we have observed its occurrence in many forms. Examples include the additivity bias, in which parts of complex entities that have been studied in isolation are assumed to retain their characteristics when the parts are

reintegrated into the whole from which they were drawn; the discreteness bias, in which continuously dimensioned attributes (like length) are bifurcated to their poles and continuous processes are instead segmented into discrete steps; and the compartmentalization bias, in which conceptual elements that are in reality highly interdependent are instead treated in isolation, missing important aspects of their interaction (see Coulson et al., 1989; Feltovich et al., 1989; Myers et al., 1990; Spiro et al., 1989 for presentations and discussion of the many reductive biases that have been identified). Of course, the employment of strategies of this kind is not a problem if the material is simple in ways consistent with the reductive bias. However, if real complexities exist and their mastery is important, such reduction is an inappropriate oversimplification.

- Errors of oversimplification can compound each other, building larger scale networks of durable and consequential misconception.
- The tendency towards oversimplification applies to all elements of the learning process, including cognitive strategies of learning and mental representation, and instructional approaches (from textbooks to teaching styles to testing). These various sources of simplification bias reinforce each other (e.g., one is more likely to oversimplify if an inappropriately easier learning strategy is also employed in textbooks or teaching because it is simple).

As we will see in the next section, more appropriate strategies for advanced learning and instruction in ill-structured domains are in many ways the opposite of what works best for introductory learning and in more well-structured domains. For example, compartmentalization of knowledge components is an effective strategy in well-structured domains, but blocks effective learning in more intertwined, ill-structured domains that require high degrees of knowledge interconnectedness. Instructional focus on general principles with wide scope of application across cases or examples works well in well-structured domains (this is one thing that makes these domains well-structured), but leads to seductive misunderstandings in ill-structured domains, where across-case variability and case-sensitive interaction of principles vitiates their force. Well-structured domains can be integrated within a single unifying representational basis, but ill-structured domains require multiple representations for full coverage. For example, consider one kind of single unifying representation, an analogy to a familiar concept or experience. We have found that a single analogy may help at early stages of learning, but actually interfere with more advanced treatments of the same concept later on (Spiro et al., 1989; see also Burstein & Adelson, 1990). Any single analogy for a complex concept will always be limited in its aptness, and misconceptions that will develop when the concept is treated more fully can be predicted by knowing the ways in which the introductory analogy is misleading about or under-represents the material to be learned. To summarize, we have found that the very things that produce initial success for the more modest goals of introductory learning may later impede the attainment of more ambitious learning objectives.

There is much that appears to be going wrong in advanced learning and instruction (see also GPEP, 1984; Perkins & Simmons, 1989). The cognitive theories and instructional practices that work well for introductory learning and in well-structured domains not only prove inadequate for later, more advanced treatments of the same topics, but adherence to those theories and practices may produce impediments to further progress. Our conclusion is that a

reconceptualization of learning and instruction is required for advanced knowledge acquisition in ill-structured domains (see also Spiro et al., 1987, 1988, 1989; Spiro & Jehng, 1990; Feltovich, Spiro, & Coulson, In press). Such a reconceptualization, taking into account the problems posed by domain ill-structuredness and the patterns of advanced learning deficiency observed in our studies, is presented next, in our discussion of constructivism and a new constructive orientation, Cognitive Flexibility Theory. After a brief survey of the tenets of that theory, we show its implications for the design of computer hypertext learning environments that are targeted to the features of difficulty faced by advanced learners in ill-structured domains.

4.3.3 Constructivism. Old and New: Cognitive Flexibility Theory and the Promotion of Advanced Knowledge Acquisition

The orientation that we have been discussing falls within the tradition of constructivism in cognitive research. Our interpretation of this term, as it is applied to learning and instruction, is complex. We argue that there are different points in cognitive acts where constructive mental processes occur. First, we take it as an accepted cognitive principle that understanding involves going beyond the presented information. For example, what is needed to comprehend a text is not solely contained in the linguistic and logical information coded in that text. Rather, comprehension involves the construction of meaning: the text is a preliminary blueprint for constructing an understanding. The information contained in the text must be combined with information outside of the text, including most prominently the prior knowledge of the learner, to form a complete and adequate representation of the text's meaning (see Spiro, 1980, for a review; also see Ausubel, 1968; Bartlett, 1932; Bransford & Johnson, 1972; and Bruner, 1963).

However, our approach to constructivist cognition goes beyond many of the key features of this generally accepted view (see Spiro et al., 1987). The interpretation of constructivism that has dominated much of cognitive and educational psychology for the last 20 years or so has frequently stressed the retrieval of organized packets of knowledge, or schemas, from memory to augment any presented information that is to be understood or any statement of a problem that is to be solved. We argue that conceptual complexities and across-case inconsistencies in ill-structured knowledge domains often render the employment of prepackaged ("precompiled") schemas inadequate and inappropriate. Rather, because knowledge will have to be used in too many different ways for them all to be anticipated in advance, emphasis must be shifted from the retrieval of intact knowledge structures to support the construction of new understandings, to the novel and situation-specific assembly of prior knowledge drawn from diverse organizational loci in preexisting mental representations. That is, instead of retrieving from memory a previously packaged "prescription" for how to think and act, one must bring together, from various knowledge sources, an appropriate ensemble of information suited to the particular understanding or problem solving needs of the situation at hand. Again, this is because many areas of knowledge have too diverse a pattern of use for single prescriptions, stored in advance, to cover enough of the cases that will need to be addressed. (For other discussions of issues related to cognitive flexibility and "inert knowledge," see Bereliter & Scardamalia, 1985; Bransford, Franks, Vye, & Sherwood, 1989; Brown, 1989; and Whitehead, 1929.)

Thus, in Cognitive Flexibility Theory, a new element of (necessarily) constructive processing is added to those already in general acceptance, an element concerned primarily with the flexible use of preexisting knowledge (and, obviously, with the acquisition and representation of knowledge in a form amenable to flexible use). (However, also see Bartlett's, 1932, notion of "turning round upon one's schema.") This "new constructivism" is doubly constructive: 1) understandings are constructed by using prior knowledge to go beyond the information given; and 2) the prior knowledge that is brought to bear is itself constructed, rather than retrieved intact from memory, on a case-by-case basis (as required by the across-case variability of ill-structured domains). (Also see Bereiter, 1985.) Cognitive Flexibility Theory is a "new constructivist" response to the difficulties of advanced knowledge acquisition in ill-structured domains. It is an integrated theory of learning, mental representation, and instruction. We now turn our attention to that theory. (Having discussed the relationship of Cognitive Flexibility Theory to constructivism, the latter term will not be used explicitly very often in the remainder of the chapter -- but it should be understood that when we talk about Cognitive Flexibility Theory, we are referring to a particular constructivist theory.)

Cognitive Flexibility Theory: A Constructivist Approach to Promoting Complex Conceptual Understanding and Adaptive Knowledge Use for Transfer

Limitations of space will not permit a detailed treatment of the key features of Cognitive Flexibility Theory in this section. Let it suffice to say that the tenets of the theory are direct responses to the special requirements for attaining advanced learning goals, given the impediments associated with ill-structured features of knowledge domains and our findings regarding specific deficiencies in advanced learning -- knowing what is going wrong provides a strong clue for how to fix it. In lieu of any comprehensive treatment, we will discuss here one central aspect of the theory. Then, we will show how that aspect creates implications for the design and use of hypertext learning environments. For more detailed treatments of Cognitive Flexibility Theory, see Spiro et al. (1987, 1988), Spiro and Jehng (1990) and Fellovich et al. (in press).

The aspect of Cognitive Flexibility Theory that we will briefly discuss here and use for illustrative purposes involves the importance of multiple juxtapositions of instructional content. Some other aspects of the theory will be referred to in passing in the context of that discussion. (Many key tenets of Cognitive Flexibility Theory will not be mentioned at all; e.g., the vital importance of students' active participation in learning). A central claim of Cognitive Flexibility Theory is that revisiting the same material, at different times, in rearranged contexts, for different purposes, and from different conceptual perspectives is essential for attaining the goals of advanced knowledge acquisition (mastery of complexity in understanding and preparation for transfer). Content must be covered more than once for full understanding because of psychological demands resulting from the complexity of case and concept entities in ill-structured domains, combined with the importance of contextually induced variability and the need for multiple knowledge representations and multiple interconnectedness of knowledge components (see Spiro et al., 1988, for justifications of all these requirements). Any single explanation of a complex concept or case will miss important knowledge facets that would be more salient in a different context or from a different intentional point of view. Some of the representational perspectives necessary for understanding will be grasped on a first or second exploration, while others will be missed until further explorations are undertaken. Some

useful connections to other instructed material will be noticed and others missed on a single pass (with connections to nonadjacently presented information particularly likely to be missed). And so on. Revisiting material in an ill-structured domain is not a simple repetitive process useful only for forming more durable memories for what one already knows. For example, re-examining a case in the context of comparison with a case different from the comparison context (i.e., the first time the case was investigated) will lead to new insights (especially if the new "reading" is appropriately guided); this is because partially nonoverlapping aspects of the case are highlighted in the two different contexts. The more complex and ill-structured the domain, the more there is to be understood for any instructional topic; and therefore, the more that is unfortunately hidden in any single pass, in any single context, for any restricted set of purposes, or from the perspective of any single conceptual model.

For example, consider the importance of multiple knowledge representations, which is one thing made possible by multiple passes through the same material. A key feature of ill-structured domains is that they embody knowledge that will have to be used in many different ways, ways that can not all be anticipated in advance. Knowledge that is complex and ill-structured has many aspects that must be mastered and many varieties of uses that it must be put to. The common denominator in the majority of advanced learning failures that we have observed is oversimplification, and one serious kind of oversimplification is looking at a concept or phenomenon or case from just one perspective. In an ill-structured domain, that single perspective will miss important aspects of conceptual understanding, may actually mislead with regard to some of the fuller aspects of understanding, and will account for too little of the variability in the way knowledge must be applied to new cases (Spiro et al., 1989). Instead, one must approach all elements of advanced learning and instruction with the tenet of multiple representations at the center of consideration.

Cognitive Flexibility Theory makes specific recommendations about multiple approaches that range from multiple organizational schemes for presenting subject matter to instruction to multiple representations of knowledge (e.g., multiple classification schemes for knowledge representation). Knowledge that will have to be used in a large number of ways has to be organized, taught, and mentally represented in many different ways. The alternative is knowledge that is usable only for situations like those of initial learning; and in an ill-structured domain that will constitute just a small portion of the situations to which the knowledge may have to be applied.

Given all of this, it should not be surprising that the main metaphor we employ in the instructional model derived from Cognitive Flexibility Theory (and in our related hypertext instructional systems) is that of the criss-crossed landscape (Spiro et al., 1987; Wittgenstein, 1953), with its suggestion of a nonlinear and multidimensional traversal of complex subject matter, returning to the same place in the conceptual landscape on different occasions, coming from different directions. Instruction prepares students for the diversity of uses of ill-structured knowledge, while also demonstrating patterns of multiple interconnectedness and context dependency of knowledge, by criss-crossing the knowledge domain in many ways (thereby also teaching students the importance of considering complex knowledge from many different intellectual perspectives, tailored to the context of its occurrence). This should instill an epistemological belief structure appropriate for ill-structured domains and provide a

repertoire of flexible knowledge representations that can be used in constructing assemblages of knowledge, taken from here and from there, to fit the diverse future cases of knowledge application in that domain.

4.3.4 Constraints on the Design of Hypertext Learning Environments Drawn from Implications of Cognitive Flexibility Theory

Thus far, we have discussed the relationship between the nature of ill-structured knowledge domains and difficulties in the attainment of advanced learning goals (mastery of complexity and transfer to new situations). A principle of Cognitive Flexibility Theory was then introduced as one antidote to the problems of advanced knowledge acquisition in ill-structured domains. Now, we will briefly point to some of the ways that these cumulative considerations impinge on the design and use of hypertext learning environments.

First, the preceding discussion should make it reasonably clear that hypertext environments are good candidates for promoting cognitive flexibility in ill-structured domains. We have referred to the need for rearranged instructional sequences, for multiple dimensions of knowledge representation, for multiple interconnections across knowledge components, and so on. Features like these correspond nicely to well known properties of hypertext systems, which facilitate flexible restructuring of instructional presentation sequences, multiple data codings, and multiple linkages among content elements. It appears straightforward that a nonlinear medium like hypertext would be very well suited for the kinds of "landscape criss-crossings" recommended by Cognitive Flexibility Theory (and needed in ill-structured knowledge domains; see also Bednar, Cunningham, Duffy, & Perry, 1991).

However, it is not that easy. Implementing Cognitive Flexibility Theory is not a simple matter of just using the power of the computer to "connect everything with everything else." There are many ways that hypertext systems can be designed, and there is good reason to believe that a large number of those do not produce successful learning outcomes (e.g., because they lead the learner to become lost in a confusing labyrinth of incidental or ad hoc connections). What is needed is the discipline of grounding hypertext design in a suitable theory of learning and instruction. That is what we have done in several prototype hypertext systems derived from Cognitive Flexibility Theory and tailored to the known obstacles to advanced learning in difficult and ill-structured domains (Spiro et al., 1988; Spiro & Jehng, 1990). To provide some idea of how theory informs design, consider just one very simple example of a hypertext design decision that responds to an aspect of Cognitive Flexibility Theory-based logic discussed in the last section: rearrangement of the presentation sequence of content (that has been investigated previously), in order to produce different understandings when that content is "re-read."

Illustrating the Theory- and Context-Based Logic of Hypertext Design

Because of the feature of conceptual instability in ill-structured domains (i.e., the same conceptual structure takes on many more meanings across instances of its use than in well-structured domains), Cognitive Flexibility Theory dictates, as discussed in the last section, that one kind of instructional revisiting should produce an appreciation in the learner of the varieties of meaning "shades" associated with the diversity of uses. As Wittgenstein argued (1953), the meaning of ill-structured concepts is in their range of uses, rather than in

generally applicable definitions--there is no simple "core meaning." We extend Wittgenstein's claim to larger units than the individual concept (e.g., complex conceptual structures such as a theme of a literary work). So, a feature built into our hypertexts is conceptual structure search: content is automatically re-edited to produce a particular kind of "criss-crossing" of the conceptual landscape that visits a large set of case examples of a given conceptual structure in use. The learner then has the option of viewing different example cases in the application of a concept he or she chooses to explore. That is, the instructional content is re-edited upon demand to present just those cases and parts of cases that illustrate a focal conceptual structure (or set of conceptual structures). Rather than having to rely on sporadic encounters with real cases that instantiate different uses of the concept, the learner sees a range of conceptual applications close together, so conceptual variability can easily be examined. Learning a complex concept from erratic exposures to complex instances, with long periods of time separating each encounter, as in natural learning from experience, is not very efficient. When ill-structuredness prevents telling in the abstract how a concept should be used in general, it becomes much more important to show together the many concrete examples of uses. In sum, a hypertext design feature is incorporated as a response to a learning difficulty caused by a characteristic of ill-structured knowledge domains. (Of course, the issue of example selection and sequencing in concept instruction has been dealt with before, e.g., Tennyson & Park, 1980. What is novel about the present approach is the particular way that this issue is addressed and the kinds of higher-order conceptual structures that are studied. Even more important is the fact that that single issue is addressed within a larger, integrative framework. That is, the treatment of conceptual variability is just one aspect of a complete approach in which the diverse aspects are theoretically united.)

Following this same kind of logic, we will sketch briefly some of the other ways that hypertext design features can be made to match the goals of advanced learning--under the constraints of domain ill-structuredness and according to the tenets of Cognitive Flexibility Theory. For this purpose we will use one of our Cognitive Flexibility Hypertext prototypes, Exploring Thematic Structure in Citizen Kane (Knowledge Acquisition in Nonlinear Environments--"KANE," for short; see Spiro and Jehng, 1990, for details), which teaches processes of behavior interpretation in a post-structuralist mode (e.g., Barthes, 1967).

KANE is a learning environment that goes beyond typical instructional approaches to literary interpretation that too often settle on a single, integrative understanding ("The theme of Citizen Kane is X."). Instead, students are shown that literary texts (in this case a videodisc of a literary film) support multiple interpretations, the interpretations combine and interact, they take on varying senses in different contexts, and so on. For example, the issue of conceptual variability that was discussed above is addressed by providing an option that causes the film to be re-edited to show just those scenes that illustrate any selected conceptual theme of the Kane character (e.g., "Wealth Corrupts"). Using this option, the learner could, for example, see five scenes in a row, taken from various places in the film, that illustrate different varieties or "flavors" of the "Wealth Corrupts" theme. Each scene essentially forms a miniature case of the Kane character's behavior that illustrates the targeted theme. (Although the student is assumed to have already seen the film one or more times--this is advanced knowledge acquisition for Citizen Kane--the nonlinear presentation may still occasionally confuse. Therefore, to deal with this and other kinds of out-of-sequence criss-crossings, a design feature of Cognitive Flexibility Hypertexts is the provision of optional background

information on the contexts immediately preceding the one being explored.) Because of the inability of abstract definitions (as might be construed for a theme such as "Soulless Man") to cover conceptual meanings-in-use in ill-structured domains, supplementary guidance about the way meaning is used in a particular situation (Brown, Collins, & Duguid, 1989) is required. This is provided for in KANE by giving the learner the option of reading an expert commentary on the special shade of meaning associated with the conceptual theme, as applied to a scene, immediately after the scene is viewed. These functional and context-sensitive (particularized) definitions explain why the scene is considered to be a case of a theme such as "Wealth Corrupts." Note that a particularized representation of meaning is not the same as a dictionary sense of a word: the latter refers to different sub-types of a word's meaning, but with an implied similarity or overlap across instances of the same type--so there is less need to tailor to the individual case; in contrast, particularizing, as we mean it, implies a representation of a concept that is necessarily expressed in terms of an instance of usage (case, example, scene, occasion of use), as required in an ill-structured domain. Commentaries also include information about knowledge access: what cues in the case context should provide a "tip-off" that a particular concept might be relevant for analyzing a case--if one can not access relevant conceptual information in memory, this knowledge will not be useful on subsequent occasions.

The commentaries also provide cross-references to other instantiations of the conceptual structure that constitute an instructionally efficacious set of comparisons (e.g., other cases/scenes in which either a roughly similar or saliently different particularized sense of that conceptual theme occurs). The guiding commentaries also include another important kind of cross-reference, namely to other conceptual themes that have interpretive relevance in accounting for the same case of Kane's behavior, concepts that interact with and influence the meaning of each other in that scene. (Note that these different kinds of cross-references counter the reductive tendencies toward compartmentalization of concepts and their cases of application that we have found to be harmful in advanced learning.) Thus there is a double particularization in Cognitive Flexibility Hypertexts: the generic conceptual structure is particularized not only to the context of a specific case, but also to the other concepts simultaneously applicable for analyzing that case. That is, each case or example is shown to be a complex entity requiring for its understanding multiple conceptual representations, with the role of non-additive conceptual interdependencies highlighted.

Each of the conceptual themes used in KANE is itself a wide-scope interpretive schema that has been argued for in the secondary literature on the film as being the most important theme for understanding the character of Kane. In reality, however, an ill-structured domain has no single schema that is likely to cover everything of interest for an individual case, nor is any schema/theme/concept likely to dominate across a wide range of cases. Therefore, the greater the number of such broad-gauge schemata that are available (and KANE provides ten), the greater the utility for understanding in two senses. First, there will be adequate coverage of the complexity of an individual case by an appropriately diverse set of schemata (something which is also modeled in KANE by the simultaneous display of all the relevant conceptual themes in each scene). Second, the likelihood is increased that the most apt set of conceptual schemata will be cognitively available for understanding any one of the highly diverse new cases that will be encountered in an ill-structured domain--the more conceptual structures there are to choose from, each a powerful schema itself and each taught in its complex diversity of patterns

of use, the greater the chance that you will find a good fit to a given case. A related virtue is that configurations of combinations of conceptual structures are thereby demonstrated; since multiple conceptual representations will be required for each instance of knowledge application, the ability to combine conceptual entities and to recognize common patterns of their combination is crucial. The process of situation-specific knowledge construction, so important for transfer in ill-structured domains, is thus supported in at least two important ways: the processes of adaptive knowledge assembly are demonstrated, and the flexible knowledge structures required for this assembly are acquired. Furthermore, as users of the program shift over time into more of a "free exploration" mode, where they independently traverse the themes of the film in trying to answer questions of interpretation (posed by teachers or themselves), their active participation in learning the processes of knowledge assembly increases.

Flexible tools for covering content diversity and for teaching knowledge assembly combine to increase the resources available for future transfer/application of knowledge (e.g., interpreting a scene that has not yet been viewed or assembling prior knowledge to facilitate comprehension of a critique written about the film). By making many potential combinations of knowledge cognitively available--either by retrieval from memory or by context-sensitive generation--the learner develops a rich palette to paint a knowledge structure well fit to helping understand and act upon a particular case at hand. This is especially important in an ill-structured domain because there will be great variety in the demands on background knowledge from case to case (and with each case individually rich in the knowledge blend required).

This discussion could continue for many other features of hypertext learning environments that are specifically derived from Cognitive Flexibility Theory. What would be in common across any such discussion is that each feature could be shown to have the following purpose: to counter an advanced learning difficulty endemic to ill-structured domains.

4.3.5 Concluding Remarks

We have just discussed a few of the many kinds of revisitings of instructional content in rearranged contexts that are implied by Cognitive Flexibility Theory and embodied in our hypertext systems. However, our goals in this chapter were necessarily limited. Our purpose was merely to begin to illustrate the way design features of a particular kind of computer learning environment are related to cognitive and instructional theories that are themselves based on the problems posed by the interaction of learning objectives and characteristics of ill-structured knowledge domains. That is, our intention was to illustrate a way of thinking about the design of hypertext learning environments that is sensitive to and dependent upon the cognitive characteristics necessary for advanced knowledge acquisition in ill-structured domains. In particular, these are the characteristics of the "new constructivism" that we discussed earlier and that are properties of Cognitive Flexibility Theory. The realm of constructive processes must be taken beyond the retrieval of knowledge structures from memory (for the purpose of "going beyond the information given" in some learning situation), to also include the independent, flexible, situation-specific assembly of the background knowledge structures themselves.

In sum, we consider our work to be moving toward a systematic theory of hypertext design to provide flexible instruction appropriate for developing cognitive flexibility. We have called the instructional theory that is derived from Cognitive Flexibility Theory and applied in flexible computer learning environments Random Access Instruction. It, and the developing hypertext theory, is laid out in considerable detail in Spiro and Jehng (1990). We are encouraged so far about the robustness, systematicity, and generality of our hypertext design principles, in that they have been applied in very similar ways to develop hypertext prototypes in domains as diverse as cardiovascular medicine, literary interpretation, and military strategy. Preliminary data on the effectiveness of these Cognitive Flexibility Hypertexts are also encouraging. For example, Jacobson and Spiro (1991a) investigated two different design approaches for structuring a hypertext learning environment to provide instruction in a complex and ill-structured domain (the social impact of technology). The results of this experiment revealed that while the design that emphasized the mastery of declarative knowledge led to higher performance on measures of memory for presented facts, the design based on Cognitive Flexibility Theory (which highlighted different facets of the material by explicitly demonstrating critical interrelationships between abstract and case-centered knowledge components, in multiple contexts on different passes through the same content) promoted superior transfer to a new problem solving situation. More empirical testing is clearly required, and numerous other issues of hypertext design remain to be discussed. However, those are stories for another time.

4.3.6 Footnote to Chapter

The main contributors to this chapter were Rand J. Spiro, Paul J. Feltovich, Michael J. Jacobson, and Richard L. Coulson.

4.4 Cognitive Flexibility, Random Access Instruction, and Hypertext: Theory and Technology for the Nonlinear and Multidimensional Traversal of Complex Subject Matter

In the preceding chapter, we provided a rationale for hypertext instruction in ill-structured domains, and we introduced our hypertext approach and sketched some of its features. In the present chapter we will present an extended and detailed statement of our novel theory of hypertext design. We will clearly demonstrate the basis in cognitive and instructional theory, derived from our own work (as discussed throughout the preceding sections of this Report) and the work of others, for the design decisions made in building what we call Cognitive Flexibility Hypertexts.

4.4.1 Introduction

Traditional methods of instruction rely upon linear media (e.g., textbooks and lectures). Linearity of media is not a problem when the subject matter being taught is well-structured and fairly simple. However, as content increases in complexity and ill-structuredness, increasingly greater amounts of important information are lost with linear approaches and the unidimensionality of organization that typically accompanies them. The advent of random access computer technologies makes practicable new forms of nonlinear and multidimensional learning and instruction that are better suited to conveying complex content. For example, it becomes a straightforward matter to revisit the same content material in a variety of different contexts, with each visit bringing out additional aspects of that content's complexity that are missed in the single pass of linear coverage. We use the expression "random access instruction" to refer to a cluster of fundamental issues brought into play by nonlinear learning with random access media. This chapter discusses a unified theoretical approach to those fundamental issues, an approach that provides a foundation of principles to guide random access instruction.

The research to be discussed in this chapter illustrates aspects of a general theoretical orientation to knowledge acquisition and application in complex content domains, Cognitive Flexibility Theory (Spiro, Coulson, Fellovich, & Anderson, 1988; Spiro, Vispoel, Schmitz, Samarapungavan, & Boerger, 1987). We will argue for the suitedness of this theory to the special needs of random access instruction. The generality of the theory is attested to by the fact that it has been applied in roughly the same way to such diverse areas as literary comprehension and interpretation, biomedical cognition, history, and military strategy (Fellovich, Spiro, & Coulson, 1989; Spiro, Fellovich, Coulson, & Anderson, 1989; Spiro et al., 1987, 1988).

Although the present chapter is primarily concerned with the application of theoretical issues of advanced learning on complex topics to a particular instructional approach, there will also be a focus on one illustrative context for that discussion: a computer program in the area of literary comprehension and the understanding of complex patterns of individual behavior. Again, though, the chapter has a general theoretical point, not limited to the one area chosen to illustrate the theoretical concerns. The interest of this domain for the military takes several forms: First, its structure is like other ill-structured domains (including the military strategy of the indirect approach; see Chapter 4.6), so the lessons learned from this *manageable*

first step provide direction for more complicated succeeding steps (e.g., that in Chapter 4.6). And the managability of this domain makes it ideal for instantiation of detailed developments in the cognitive and instructional theory, as will be seen throughout this chapter. Second, the program makes use of *videodisc*, a technology of great interest to the military. Third, predicting and explaining complex behavior has some direct military relevance (e.g., interpretation of an opposing commander's personality as it relates to his actions, which is a very difficult thing to train in *military intelligence*).

Because the meaning of complex real-world situations, as well as literary works, transcends the mere chronicling of the events they contain and also tends not to be reducible to any single interpretation, representation of patterns of thematic deployment often form the most useful scaffolding for their comprehension (rather than a single event-oriented schema). For example, a literary work usually contains many themes and symbolic structures. A given theme is likely to be used in diverse ways, be relevant at irregular intervals throughout the work, and form intricate patterns of combination with the other themes. We argue that mastery of complexities such as these, which are found in all rich content domains, is greatly facilitated by the kind of guided nonlinear and multidimensional explorations typified by the instructional program (based on Cognitive Flexibility Theory) discussed in this chapter. (Furthermore, we have accumulated considerable evidence that more advanced understandings of this kind tend very frequently not to be successfully attained in typical instructional settings--see, for example, Feltovich et al., 1989 and Spiro et al., 1988, 1989.)

The project that is the primary focus of this chapter uses a random access (i.e., interactive, programable) videodisc of the complexly structured classic film Citizen Kane (Welles & Stein, 1984). A large number of short segments drawn from several scenes of the film have been closely analyzed for their patterns of instantiation of the film's thematic and symbolic structure. This analysis forms the basis for a very large number of computer-generated re-editings of the segments as a function of their patterns of thematic relatedness. In one use of the program, after students/subjects see the full scenes in their natural sequence, they are then exposed to variously rearranged traversals of the scenes in different contextual arrangements. Commentary is appended explicating the thematic and symbolic contrasts illustrated by each traversal (i.e., each thematically based re-editing of the scenes).

This use of programable video to construct multiple 'texts' is intended to prepare learners (better than instruction with a standard linear schema for the film would) to go beyond the ability to merely reproduce the instruction they received, and instead to be able to independently apply the instructed knowledge to new situations that differ in their characteristics from those of initial learning. That is, the goal is transfer. An example of a transfer goal that the program is intended to facilitate is an improved ability to comprehend scenes from the film that have not yet been seen. Or the knowledge acquired from instruction should serve more effectively as background knowledge to support the comprehension of critical texts written about the film. It is also hoped that more general transfer might occur: by having the need for complex analysis of literary works clearly demonstrated to students, they will be more likely to avoid typical tendencies toward oversimplification in other literary works; and by fostering skill in the processes of complex text-based interpretation, the students should be better able to carry out the more complex analyses of new works. More generally still, this

instructional program could be combined with similar ones developed in other domains to form a staging ground to alert students to generic hazards of oversimplification and to prepare them to deal with complexity across domains. (Note that an aspect of the instructional approach that is considered crucial for transfer but that receives less attention here is the active involvement of the learner. The role of student participation and exploration becomes clearer in the detailed example presented later in the chapter.)

4.4.2 Plan of the Chapter and Overview

The chapter begins by posing a set of problems for cognitively based theories of learning and instruction. These problems revolve around a concern for advanced knowledge acquisition (i.e., post-introductory learning) in a content area. At advanced stages of knowledge acquisition content becomes more complex and the relationships across the cases that knowledge has to be applied to become more irregular. We call domains that have these features of content complexity and irregularity of application contexts ill-structured domains. At the same time that greater ill-structuredness must be dealt with by advanced learners, the goals of learning shift: (a) from the attainment of superficial familiarity with concepts and facts to the mastery of important aspects of conceptual complexity, and (b) from knowledge reproduction to knowledge use (transfer, application). (See Feltoovich et al., 1989 and Spiro, et al., 1988, 1989, for detailed discussions of advanced knowledge acquisition.)

After presenting a brief discussion of problems in dealing with advanced knowledge acquisition in ill-structured domains, we present a remedy: the "Cognitive Flexibility Theory" of learning, knowledge representation, and knowledge transfer. By "cognitive flexibility" we mean the ability to spontaneously restructure one's knowledge, in many ways, in adaptive response to radically changing situational demands (both within and across knowledge application situations). This is a function of both the way knowledge is represented (e.g., along multiple rather than single conceptual dimensions) and the processes that operate on those mental representations (e.g., processes of schema assembly rather than intact schema retrieval).

We then discuss the instructional approach systematically derived from that theory and implemented in computer learning environments, the "Cognitive Flexibility Hypertext" approach to random access instruction. Key features of the approach are illustrated by a functioning hypertext prototype that teaches aspects of the complex multithematic structure of the film Citizen Kane using random access videodisc. The generality of the approach will be underscored by a brief presentation of highlights of Cognitive Flexibility Hypertext prototypes in other domains (e.g., cardiovascular medicine and military strategy).¹

We will be especially concerned with how the knowledge representations that are built by using programs like KANE support the application of old knowledge in new situations. We stress the process of situation-dependent schema assembly as against the views of schema theories that depend upon the mere retrieval of precompiled generic knowledge structures that are monolithically superimposed on the concrete case at hand. The latter approach is simply too unwieldy to support transfer when there is a lot of complexity in the individual case of application and when cases vary considerably, one to the next. If you may have to use knowledge

in many ways, in a diverse set of circumstances, you cannot rely on a small number of rigidly prepackaged schemas (see Schank, 1982, as well).

A note on hypertext. The term hypertext refers to computer-based texts that are read in a nonlinear fashion and which are organized on multiple dimensions (see Marchionini, 1988, for a review). The same material (which can be any kind of randomly accessible medium, e.g., text, video, audio) is capable of being explored in different ways, with the different exploration paths producing what are essentially multiple texts for the same topic. We will have much more to say about hypertext later in the chapter--in particular about how the features of our approach avoid problems commonly found in other hypertext systems, such as the problem of pre-stored links and the problem of getting lost in a labyrinth of connections between ideas. For now, however, it is worth pointing out that the outline of the chapter just presented also corresponds to a systematic response to widespread problems in the development of hypertext learning systems. In particular, work on hypertext has tended to be atheoretical, driven by the power of the technology rather than by a clear sense of how to respond to two key issues: (a) the stages of learning and purposes of learning for which this unfamiliar kind of instructional environment may be best suited--not all kinds of learning require so complex and potentially confusing an approach; and (b) the cognitive psychology of nonlinear learning--hypertext systems would be easier to use and would support greater educational attainment if they were systematically designed in accordance with a theory of how the information will be processed, mentally represented, and later used (see Spiro, et al., 1988). This chapter can be thought of as being sequenced to correspond to these questions about hypertext (as well as forming a natural sequence for the presentation of problems of advanced learning and their solutions--as described in the overview).²

Our answer to the questions that we said should be posed for any hypertext systems are that they are best suited for advanced learning, for transfer/application learning goals requiring cognitive flexibility, in complex and ill-structured domains--rather than introductory learning, for memory tests, in simpler domains. [Incidentally, we were able to avoid the endemic atheoretical approaches to hypertext development only because of a fortunate coincidence of sorts: our work on the cognitive psychology of learning for transfer in ill-structured domains happened to precede our interest in computer-supported instruction (Spiro, 1980; Spiro & Myers, 1984; Spiro et al., 1987).]

4.4.3 Advanced Knowledge Acquisition and the Problems of Ill-Structuredness, Cognitive Flexibility, and Knowledge Transfer

Two important things happen as you move beyond the initial introduction to a content area to more advanced stages of knowledge acquisition in that area: (1) the conceptual content tends to become more complex and the basis of its application more ill-structured; and (2) the goals of learning and the criteria by which learning is assessed shift (or should shift): (a) from superficial or introductory level familiarity with concepts to the mastery of important aspects of complexity (despite their difficulty); and (b) from accurate reproductive memory and imitative rule-following for instructed material to the ability to apply what was taught in new and greatly varying contexts--knowledge transfer.³

What do you have to do to attain a deep understanding of a complex concept, to "get it right"?²⁴ What has to be done to be ready to apply conceptual knowledge in a domain where the phenomena occur in irregular patterns? How can one be prepared to use knowledge in the great variety of ways that may be required in a rich domain? Unfortunately, the indications are that what should be done in all of these cases is very often just the opposite of what tends to be done in conventional instructional practices (Spiro et al., 1987). Learning and instruction for mastery of complexity and application in a complex and ill-structured domain cannot be compartmentalized, linear, uniperspectival, neatly hierarchical, simply analogical, rigidly prepackaged (Spiro et al., 1987, 1988). Yet it much too often is, and the result is the development of widespread and serious misconceptions and difficulties in knowledge application (Feltovich et al., 1989; Spiro et al., 1988, 1989). This is because complex and ill-structured knowledge domains are characterized by such features as nonuniformity of explanation across the range of phenomena to be covered, nonlinearity of explanation, nonadditivity following decomposition, context-dependency, irregularity of overlap patterns across cases (reducing the effectiveness of prototypes and simple analogies), absence of wide scope defining features for category application, and so on (see Spiro et al., 1987, 1988, for discussions of domain ill-structuredness).

We have shown in a number of studies that the learning of complex content material in ill-structured domains requires multiple representations--multiple explanations, multiple analogies, multiple dimensions of analysis (see, for example, Spiro et al., 1987; also see White & Frederiksen, 1987). Mental representations need to be open rather than rigid and closed; nonlinear instructional sequences need to be followed to avoid missing key points; assumptions of regularity and homogeneity have to be replaced by acknowledgement of irregularity and heterogeneity. Learning that has these characteristics of openness and plurality produces cognitive flexibility: the ability to adaptively re-assemble diverse elements of knowledge to fit the particular needs of a given understanding or problem solving situation. In an ill-structured domain, one cannot fit the wide variety of real-world cases of a given type that will be encountered to the same "plaster-cast" knowledge structure (although a common failing of advanced learners is that they will try very hard to do this).

Standard technologies (e.g., books, lectures, etc.) are not well suited to these requirements for the development of cognitive flexibility. But the problem goes beyond the limitations of traditional instructional technologies: theories of cognition and instruction too often focus either on introductory learning or advanced learning in well-structured domains. They therefore have features that are antithetical to those required to deal with the complexities that need to be mastered at advanced stages of knowledge acquisition, as described above. (See Spiro & Myers, 1984 and Spiro et al., 1987 for discussions of limitations of earlier versions of schema/frame/script theories. Cognitive Flexibility Theory is a response to those limitations.)

Actually, although there has been considerable research attention devoted to differences between experts and novices, the intermediate stage of advanced knowledge acquisition that bridges between novicehood and expertise remains little studied. This neglect has had serious consequences given the indications from our research that characteristics of early learning are inimical to advanced knowledge acquisition--many of the strategies of learning and instruction that are most successful in introductory learning (e.g., the use of analogy) form impediments to the eventual development of more sophisticated understandings (Feltovich et al., 1989; Spiro et al., 1988, 1989). So, new theories are needed, as well as new technologies appropriate to

those theories. In the next section we first discuss Cognitive Flexibility Theory and then the technological instructional orientation derived from that theory, the Cognitive Flexibility Hypertext approach.

4.4.4 Learning and Instruction for Cognitive Flexibility: Criss-Crossing Conceptual Landscapes

The central metaphor of Cognitive Flexibility Theory is the "criss-crossed landscape." The metaphor derives from Wittgenstein (1953) who, in his preface to the Philosophical Investigations, despaired that all of his attempts to weld his complex ideas into a conventionally unified exposition, to force his ideas in any single direction, crippled those ideas. Rather than reducing the complexity of his ideas for purposes of expositional elegance and (spurious) theoretical parsimony, he opted instead to write a different kind of book. He would treat the philosophical topics that were his subject as forming a complex landscape, and he would sketch those topics as sites within the landscape. He would then arrange these sketches of local regions of the landscape to form a kind of album. The sequences in the "album" would represent different traversals of the (conceptual) landscape. So in order to insure that the complex landscape would not be oversimplified, he would endeavor to "criss-cross" it in many directions; that is, the same sketches of specific issues (or cases) would reappear in different contexts, analyzed from different perspectives.

Although Wittgenstein did not explicitly make the following claim, it seems likely he would agree (and in any case we argue) that since the complexity of a single region (issue, example, case) in a landscape would not be fully graspable in any single context, its full multifacetedness would be brought out by rearranging the sequence of sketch presentations in the album so that that region would be revisited from a variety of vantage points, each perspective highlighting aspects of the region in a somewhat different way than the other perspectives. A synoptic view of the complexity of the conceptual landscape would cumulatively emerge over a number of traversals--the richness of the subject matter would not be crippled if the content was examined in many different ways.

Cognitive Flexibility Theory generalizes Wittgenstein's metaphor of the criss-crossed landscape to apply to any complex and ill-structured knowledge domain. Furthermore, the metaphor is extended beyond Wittgenstein's concern for exposition (i.e., a style of writing). We use the metaphor to form the basis of a general theory of learning, instruction, and knowledge representation. One learns by criss-crossing conceptual landscapes; instruction involves the provision of learning materials that channel multidimensional landscape explorations under the active initiative of the learner (as well as providing expert guidance and commentary to help the learner to derive maximum benefit from his or her explorations); and knowledge representations reflect the criss-crossing that occurred during learning.

By criss-crossing topical/conceptual landscapes, highly interconnected, web-like knowledge structures are built that permit greater flexibility in the ways that knowledge can potentially be assembled for use in comprehension or problem solving. The likelihood that a highly adaptive schema can be assembled to fit the particular requirements for understanding or acting in the situation at hand is increased. In other words, the range of differing situations that the knowledge could be transferred to is increased. In ill-structured knowledge domains, with

their great heterogeneity across potential instances of knowledge application, this flexibility is essential. Since one cannot have a prepackaged knowledge structure for every situation that might be encountered, the emphasis must shift from intact schema retrieval to flexibility of situation-specific schema assembly (see Spiro et al., 1987, 1988). By criss-crossing a conceptual landscape in many directions, knowledge that will have to be used in many ways is taught in many ways.

4.4.4.1 Random Access Instruction as the Ideal Medium for Criss-Crossing Ill-Structured Domains: New Kinds of Learning and Instruction for Cognitive Flexibility Made Possible by Computers

Clearly, Cognitive Flexibility Theory, with its emphasis on repeated presentations of the same material in rearranged instructional sequences and from different conceptual perspectives, is most efficiently implemented in delivery systems with random access capabilities (e.g., programable videodisc, as in the program discussed in this chapter). The "random access instruction" implementation of Cognitive Flexibility Theory acts as an antidote to the various forms of failure in advanced knowledge acquisition frequently associated with traditional learning and instruction (Coulson et al., 1989; Feltovich et al., 1989; Spiro et al., 1988). If typical approaches have overlinearized, one can construct nonlinear presentations. If material has been presented from just one point of view, one can re-present it from different points of view. Cases that have been slotted in a rigid hierarchical structure can be repeatedly re-presented to attain heterarchical or montage-like structural representations. If partially overlapping exemplars have been indiscriminately lumped under one category in a way that causes important differences among them to be missed, one can demonstrate the diversity amongst the similarity. Or, if exemplars that partially overlap in important ways have been separated into different conceptual categories, the similarity amongst the diversity can be demonstrated. If aspects of knowledge have been overly compartmentalized, their insularity can be overcome by joint presentation. If decomposed elements are not additively assemblable, they can be reassembled with a more complex combinatory logic. If an old example/case is employed too monolithically as a precedent for a new one, you can decompose examples and then recombine aspects of different examples to achieve the most accurate (the most closely fitting) set of multiple precedents for understanding in a new situation.

We call the computer-based instructional programs derived from Cognitive Flexibility Theory and built to carry out such operations as those just listed Cognitive Flexibility Hypertexts.

A note on what kind of hypertext we are talking about. The Citizen Kane program and the other Cognitive Flexibility Hypertexts have a specialized function--they are concerned with structural characteristics of the cases they try to provide help in understanding. Rather than trying to provide a diverse and entertaining experience, they are trying to simulate experience, insofar as structural aspects of experience are relevant to the knowledge representations that must be developed to support cognitive processing of subsequently encountered cases. The presented material is only that which is needed to form a fairly full representation of a case's internal structure. We are not saying that other material, such as information about an author (or the director of Citizen Kane) would not be helpful as well. It is simply our intention to deal

with one important and difficult aspect of the hypertext problem: having the hypertext convey complex knowledge structural characteristics.

A note on the use of the terms "random" and "flexible." Two potential misunderstandings should be mentioned. First, the designation "random" is not intended to denote an absence of systematic structure underlying the instructional approach. Far from it: there is an abundance of structure in Cognitive Flexibility Hypertexts, albeit a highly multidimensional and weblike structuration. The randomness concerns the computer system's ability to utilize the underlying structure in a great variety of ways and without inconvenient delays (as when you need to fast-forward videotape to get to a distant segment). The same information is easily accessed in various combinations with other, perhaps distant information in a large number of different contexts.

Second, by promoting cognitive flexibility we are not advocating an "anything goes" mentality. To say, for example, that a text may support multiple alternative interpretations is not to say that interpretations may be offered without warrant. Learners must always be encouraged to find evidence for their claims; a strength of our systems is that their structure and operation make the search for evidence more manageable.

Beyond access: A note on the need for grounding of hypertext development in a theory of learning and instruction. We refer to the systems that we have built as "hypertexts" because of their nonlinear and multidimensional nature. (See Marchionini, 1988, for a review of hypertext work.) However, we wish to emphasize again (see the earlier "note on hypertext") our belief that the debate on the hypertext notion has been too narrowly concerned with issues such as data access and too little concerned with the atheoretical character of much of the hypertext work (Spiro et al., 1988). We believe that it is vitally important that hypertext development not be divorced from underlying theories of cognition and instruction. We must know more about how people think and learn in the radically novel form required by hypertexts. We offer Cognitive Flexibility Theory and its extension to random access instruction as a grounding for hypertext approaches. Furthermore, we do not believe the additional cognitive load placed on learners by nonlinear instruction is always desirable. In more well-structured and simple knowledge domains, and, perhaps, in some introductory learning, the disadvantages of hypertext approaches may outweigh their advantages, and traditional approaches are likely to be more efficient and effective. We contend that hypertexts should be used primarily in those situations where traditional approaches would interfere with the goals of knowledge acquisition, namely, for advanced learners striving to master complexity and prepare for transfer in ill-structured knowledge domains.

4.4.5 A Cognitive Flexibility Hypertext for Multithematic Explorations of an Ill-Structured Domain: An Example of Random Access Instruction

The Cognitive Flexibility Hypertext, "Exploring Thematic Structure in Citizen Kane"⁵ (referred to by the shorthand designation "KANE," which can also be read to stand for "Knowledge Acquisition by Nonlinear Exploration"), is a random access instruction approach to

advanced understanding of the behavior of the central character in the film, behavior that has a complex and subtle structure. Segments of the film, combined with text, are systematically represented at different times, in different content combinations, in different sequences. The result of these rearrangements is multiple 'texts' about the film, for different learning purposes, produced by automatic computer re-editing. The program has been designed for use by advanced high school students and college students. The intent of the program is to go beyond traditional instructional approaches that tend to be overly linear, one-dimensional in their abstractions, and, in general, reductive of the complexity found in literary works. In addition to the goal of fostering advanced understandings of this specific film, the more general intention is to demonstrate to students the complex nature of literary comprehension and to help students to begin to build a more adequate repertoire of cognitive skills for the processing of complexity and for the application/transfer of complex knowledge to new situations. The work on KANE is part of a larger research program whose goal, currently being explored in several domains, is to discover general principles of effective advanced learning and instruction for complex subject matter.⁶ [In addition to their intended use in advanced knowledge acquisition, the hypertext programs can also function as helpful research tools for expert specialists, and of course, in experimental studies of nonlinear cognition and instruction.]

At this time, the KANE program does not try to teach all aspects of the film. Rather, our more restricted current focus is a case-based understanding of the internal "semantic" structure of the film, particularly as it bears on the Kane character (i.e., his motivations, values, beliefs, and so on). (Obviously, if we wanted to be more ambitious and teach the entire film, rather than just the complex makeup of its main character, the film's convoluted narrative style and many technical innovations would be ideally suited to the random access approach.) Also, at this time the program uses a preestablished underlying structure, albeit a highly complex, pluralistic, and flexible one. Upcoming versions will incorporate more options for users to generate their own structural schemes in addition to those provided for them. However, the current version does allow users considerable leeway in the secondary structures that they may add to the provided underlying structure; and freedom of student-initiated exploration of the landscape as it is complexly mapped by the provided structure is a key feature of the program (as its title indicates). As will be illustrated later, the provided underlying structure retains the potential for quite considerable openness in the interpretation of the work (Barthes, 1967).

"Exploring Thematic Structure in Citizen Kane": A brief description of the operation of the program

The following is a brief description of the program intended to provide just enough familiarity with its operation to form a context for the discussion of the theoretical implications of the program's features that comes after this context-setting section. It is important to remember that the main intent of this chapter is to present our theoretical perspective and the systematic approach to computer-supported instruction that follows from the theory; discussion of the KANE program per se is not the purpose of this chapter. Accordingly, the reader interested in more information about KANE is directed to a more detailed, archival description of the program that will appear in another paper.

Before going any further, however, an important disclaimer should be noted: Given the limitations of space, the accumulation of a critical mass of highly interconnected knowledge by a large number of "landscape criss-crosses" cannot be fully illustrated here. The examples that will be discussed can give only the faintest hint of the experiences provided by more extended exposure.

Preliminaries. It should be kept in mind that KANE is an advanced knowledge acquisition environment, not a tool for introductory learning. Thus, in the initial phase of working with the program, it is expected that students (college or advanced high school) will have already watched the film in its natural sequence at least once and preferably two or more times to become very familiar with it, and that they will have been thoroughly introduced to the major themes--before you can "play" with content, examining it in a variety of new and interesting ways, the content to be played with must be well learned.

Thematic organization: Multiple wide-scope "schemas of the whole". For the main instructional part of the program, an extensive and detailed thematic analysis was carried out of Side Two of the Citizen Kane videodisc, containing several scenes and lasting approximately 30 minutes. The analysis was based on ten themes chosen for their prominence in the critical literature on the film. The themes have wide scope: each theme has been put forward by some experts as being capable of providing a complete account of the Kane character's behavior, motivations, failings, and so on. In a sense, each of these complex and powerful themes has been thought of by some subset of expert analysts as providing a complete schema for understanding Kane. Examples of themes include: "Hollow Man" (Kane's inherent soullessness), "Wealth Corrupts," and "Outsized Ambition." Additionally, two of the themes involve Kane's relationships to important symbolic characters in the film, Leland and Thatcher. The course of Kane's development can be conceptualized as a dynamic movement towards and away from a complex of features crystallized in those characters (roughly: populism, friendship, and principled behavior for Leland, and wealth and emotional detachment for Thatcher).

As will be seen later, the use of several comprehensive themes makes it more likely that an appropriate one (or subset combination) will be available to be adaptively fit to the particular needs of situations encountered in the future. Furthermore, the variability of usage of each theme is stressed in instruction (see the discussion of variability in concept application). So, knowledge transfer is facilitated by having a larger number of wide-scope interpretive schemas available and by enabling learners to use each of those schemas in a flexible manner. The use of multiple organizing perspectives/schemas in the instructional program is an illustration of one of the most important recommendations of Cognitive Flexibility Theory: use multiple representations for advanced knowledge acquisition in ill-structured domains.

The choice of themes does not involve an especially delicate selection process. In an ill-structured domain, more schemas are better than less; so each time you add another credible schema you are adding to the scope of coverage. (Again, each of the themes could be treated as a full schema. The themes will also be called concepts at times. See footnote 4 for a discussion of the terminology used in this chapter.) Therefore, there is no need to adjudicate between alternative wide-scope schemas offered by different experts, discarding some because they are not as "good" as others. In an ill-structured domain any widely supported candidate schema is

likely to be useful on many occasions and less relevant on others--they are all "correct," but only to a limited extent. The more of them you have, the more likely you will be to have an available subset that is especially useful for the needs of processing some new case (keeping in mind that in ill-structured domains there is considerable variability across cases and each case is individually complex). By the time you get up to 10 or so wide-scope schemas, a lot of territory has been covered (certainly much more than with just one schema). So, given the need for an expansive theme selection policy (the theory requires that as many credible perspectives as possible be incorporated up to limits of cognitive tracking capacity, to support future knowledge transfer), there is little pressure to choose between competing themes for inclusion in the instructional program. In any case, any theme list will be insufficient by itself given the necessity of knowledge assembly in the context of new cases in ill-structured domains (i.e., new knowledge structures will typically have to be built for a new case by combining themes).

It should also be noted that it is all right for the themes to overlap somewhat (it would be difficult to find any set of themes that are mutually independent, yet still have wide scope of application). However, the themes in KANE only partially overlap: each makes some novel contribution not made by the others. (Obviously, it would be undesirable to use a set made up of very similar themes, themes that are not much more than paraphrases of each other.)

Case organization: Division into mini-case scenes. The second side of the videodisc was divided into twenty-five natural units. As will be seen later, the precise subdivision chosen is not particularly significant so long as it is a defensibly reasonable one, since any of a number of possible subdivisions produce roughly equivalent effects--the important thing is that the chosen segments provide a rich staging ground for instruction. These scenes function as mini-cases: self-contained units that are short enough to permit rapid study (they mainly range from 30-90 seconds), but rich enough to allow for the complex interplay of multiple themes. Each of these scene "units" is coded with a vector which specifies which of the eight themes and two symbolic perspectives has a relevant role in a given scene; this information is the basis for the computer program's theme-based search (see below).

An example of a mini-case scene, which will be referred to in later sections of the chapter, is the following segment of the film (that takes less than a minute to view):

In this scene, Kane has just recently taken over the newspaper "The New York Enquirer." The first edition of the newspaper since he took control of it will come out in the morning. Kane, his associates, and the newspaper staff have worked all night remaking the paper. Dawn is breaking. Kane's two exhausted associates are remarking on what a long day it has been. "A wasted day", says Kane. He complains that all they did was alter the surface of the paper. He wants to do something that is not superficial. He wants to make the Enquirer "as important to New York as the gas in this light." He then blows out the flame on the light, and the scene ends.

This is a very short scene, but it can teach many important lessons (and the number of lessons it can teach increases when the scene is contrasted with other scenes to draw additional lessons beyond those of the scene in isolation). Some of the lessons taught in the context of this 45-

second scene are illustrated by the following optional theme commentaries drawn from the program (with theme and sub-theme information included in brackets):

It's been a long day and night, but there is a weary sense of accomplishment in the group. Kane, however considers it a "wasted day." He has completely remade "The Enquirer," from top to bottom, but that is not enough for him. The changes are too superficial. He wants his newspaper to become of central importance to New York. He wants the people of the city to depend on the "Enquirer" as much as they depend on the gas for their lights. [An illustration of the Outsized Ambition theme, and its Grandeur/Sweep and Egomania sub-themes.]

By getting the people of New York to consider his newspaper to be as vital to them "as the gas in this lamp," he wants them to depend on him (through his newspaper, which Kane always treated as an extension of himself). [An illustration of the Power theme ("Control Others by their Dependence on You" sub-theme).

His reference to the gas in his lamp as typifying something that is essential to the city is ironic. Soon the gas lamp will be obsolete (as will Kane). This is one of the many instances of conspicuous misjudgment on Kane's part (as when he guarantees in 1938 that there will be no war in Europe). He underestimates the forces of change and consequently is a frequent misreader of the future (whether it is about gas as a source of power or newspapers as the dominant media for influencing people). He is constantly shown to be fallible, a trait that undermines his egomaniacally pitched ambition. [An illustration of the Fallibility/Mortality sub-theme of the Outsized Ambition theme.]

Everyone has worked all day and all night, but Kane is the only one who is still going. This is typical of the energy and enthusiasm of his early years. (His youthful energy contrasts starkly with the deathlike torpor and stagnancy of Kane when he is older.) [A negative example of the Hollow Man theme illustrating the Lifelessness sub-theme. His energy and vitality in this scene are opposed to his later lifelessness and lack of inner spark.]

A lot of the complexity of the domain has begun to be illustrated from just one mini-case lasting less than one minute. As we will see later, there are several purposes served by this design strategy of structuring in small segments: (e.g., acceleration of experience acquisition, making complexity tractable for the learner, facilitation of subsequent restructuring of knowledge, and others).

An important note that will help to avoid confusion in the exposition throughout the remainder of the chapter is that instances from the film are treated as "cases" from Kane's life. So, when our theoretical discussions refer generally to the role of cases, that means larger scenes from the film, while mini-cases correspond to very short scenes from the film. Again, scenes and mini-cases are used as interchangeable terms in this chapter.

Contextual support, thematic commentary and guidance. After a scene has been viewed the frame is frozen and several options are presented. Because the scenes are studied out of sequence, the user always has an option to request "stage-setters," information about the

context in which the scene is occurring, what has just happened before the scene, and so on. Thus, the "out of the blue" effect of nonsequential presentations is lessened. (It should be remembered, however, that the film has already been viewed in its natural sequence prior to this advanced stage of learning.) Also, of course, it is possible to use the program to watch clusters of scenes sequentially; that is, the student can study cases with a larger "grain size" than the mini-cases. This is recommended only for later stages of using of the program, in order to avoid the confusion likely to result from starting right in with a complex multithematic analysis of too large an amount of case information. [See the later section on "experience consolidation" for a discussion of the cognitive advantages of introducing complexity in "bite-size chunks."]

Each scene is presented with a text overlay in a corner of the screen listing the themes the scene contains. The themes that were targeted in the menu selection (see the section immediately below) are presented at the top of the list in a different color than the themes present in that unit but not targeted for search by the student. The student can choose to see text commentary on the particular nature of any theme's instantiation in that scene. Two important kinds of information included in the commentaries concern access/tailoring (why the particulars of that scene constitute grounds for saying that the given theme is illustrated there--how the generic theme is tailored to apply to that particular context) and across-scene relations (how a particular instantiation of a theme relates to instantiations elsewhere in the film).

Finally, the student can choose to see the scene over again, exit back to the menu, or continue with the next scene in the stack. Also, the student can examine a "road-map" of his prior explorations.

In sum, instruction on complex subject matter must be made as tractable for the student as possible. We have found in our research that ignoring complexity leads to unacceptable learning outcomes (Coulson et al., 1989; Feltoovich et al., 1989; Spiro et al., 1989). Essential content complexity and application irregularities must be faced, even if they are difficult to learn and teach. However, every effort must be made to help the student to manage that important complexity.

Theme-based exploration: Re-editing the film as a function of thematic content. Students operating in the program's self-directed mode (there is also a sequential and an experimenter-controlled mode of operation) are presented with a menu of themes and instructions for combining themes. They can choose to examine the occurrences of a single theme (which may be specified as to whether it is a positive or negative instance of a theme, if desired--for example, some event related to poverty would be a negative instance of the wealth theme). Or they can choose to search for scenes that illustrate combinations of several themes (conjunctively or disjunctively). The menu selection leads to a search for minicases/scenes whose vectors have the targeted thematic properties. For example, the student may choose to look for scenes that illustrate both the Wealth Corrupts and Hollow Man themes. Those scenes that fit the request are then put in a stack of scenes that the student can view.

In other words, in a very short amount of time, a re-edited version of the film that highlights just the theme(s) that the student wants to explore is ready to be presented. Because of the large number of scenes and themes, there are literally thousands of potential re-editings,

each of which has some instructional significance (some, obviously, having more significance than others). This plentitude of re-editings relates to one of the important lessons of ill-structured domains: because of domain irregularities and novelty resulting from the exponential explosion of the multiple factors in a complex concept, you always gain from additional experience--for example, a physician can never have "enough" experience.

Non-preprogrammed special initiatives: User-construction of interpretive essays.

Adding to the flexibility of the program is a miniature author-aiding system of sorts in which "special initiatives" to teach structures in the film that are not included in the theme menu can be developed (either by the student or by an experimenter or teacher). For example, if the instructor wanted to construct a visual essay about Oedipal themes in the film, he or she would enter a listing of the scenes that are pertinent to Oedipal interpretations. The film would then be rapidly re-edited to show the relevant scenes upon a request to study the Oedipal interpretation. Or the symbolic significance of snow could be easily demonstrated by using the special initiative option to direct a re-editing to show the several scenes from different parts of the film that utilize snow symbolically. These special initiatives can be saved and incorporated into lessons under the instructor's control.

Enabling KANE (pardon the pun): The virtues of Handy.

KANE was programed using the experimental Handy authoring language, developed by Don Nix. It would have been difficult if not impossible to incorporate all the features of KANE without Handy. And those features that could have been incorporated with other authoring aids would not have been so convenient to incorporate. Handy is a very useful tool for multi-media instructional development.

4.4.6 Critical Features in the Application of Cognitive Flexibility Theory to Random Access Instruction: Theoretical Rationale for Design Decisions in "Cognitive Flexibility Hypertexts"

In the last section, we offered a moderately uninterpreted description of the operation of KANE. With that description as background, we will now highlight a number of fundamental design decisions that were made in that instructional program (and in each of the other Cognitive Flexibility Hypertexts), calling special attention to the theoretical motivation for these decisions in light of the goals of advanced knowledge acquisition, as the achievement of those goals is constrained by the presence of complexity and ill-structuredness in a knowledge domain.

Cognitive Flexibility Hypertexts, including KANE, are the result of a large number of design decisions. Each of these decisions is motivated by Cognitive Flexibility Theory, and most of them have more than one justification. To illustrate the web of support underlying the design decisions, we will discuss in our limited space just one of them in fairly elaborate detail: the decision to use the mini-case (in KANE, a short scene) as the fundamental unit of organization for the search and presentation-display processes in the hypertexts. We will then more briefly highlight several other important design features of Cognitive Flexibility Hypertexts, with particular reference to their instantiation in KANE. The discussion will again focus on those features of learning and instruction that are affected by characteristics of complexity and

ill-structuredness, as those characteristics relate to the attainment of the goals of advanced knowledge acquisition (mastery of complexity and knowledge transferability/applicability).⁷ In other words, wherever possible, the rhetorical form for these discussions will follow the basic frame that was originally used to move from Cognitive Flexibility Theory to the instructional implementation of the theory's implications for application:

Because ill-structured domains have property X, the structural and processing requirements for attaining transfer are therefore Y. So, to get a learner to have the skills and knowledge to be able to achieve transfer, instructional systems ought to have design feature Z.

A sample instantiation of the above frame would be:

Because concepts in ill-structured domains have the property of substantial variability across their case applications, the structural and processing requirements for attaining transfer are therefore to understand the nature and scope of that variability, as well as the way that concepts get tailored to cases and cases signal the need for accessing concepts. So, to get a learner to have the skills and knowledge to be able to achieve transfer, an instructional system ought to have the following design feature (among others to address this need): the system should allow the student to sequentially study just those parts of cases that contain examples of uses of the concept; i.e., teach conceptual variability by demonstrating it in one place (accompanied by appropriate commentary on the nature of the conceptual variability).

4.4.6.1 The Mini-Case as the Fundamental Unit of Instruction: Structuring in Small Segments for Tractably Accelerated Acquisition of Case Experience and the Development of Flexibly Assemblable Knowledge

The mini-case (a segment drawn from a larger case) is the starting point for all instruction in Cognitive Flexibility Hypertexts. The rationale for this design decision is discussed in detail in this section.

Case-based exploration: The central importance of actual occurrences in ill-structured domains. In an ill-structured domain, across-case variability is, by definition, too great to allow abstract conceptual knowledge to have a dominant role. By the definition of ill-structuredness, any abstraction or generalization will inadequately account for what happens across the range of cases that the knowledge will have to be applied to. Thus, in any ill-structured domain a more case-centered approach is needed.

Cases cause problems too: The rationale for the mini-case. However, even the individual case is too complex and unwieldy a unit. Each case means too many things to be useful when treated as a monolithic unit. For example, isolated parts of a case may have instructional value but go unnoticed because they appear to lack significance within the frame of reference of the case as a whole. More seriously, intact cases used in instruction are frequently the cause of a problem in subsequent knowledge application: the tendency to map a new case completely to a single well-learned prototype, when part of the old case is misleadingly related to the new one-- in ill-structured domains multiple representations are the rule, and that includes understanding a new case by reference to multiple prototypes. When cases are taught monolithically, their representations are harder to pull apart in the way that is necessary for

transfer-enabling flexible reorganization of knowledge. Thus, in an ill-structured domain, an intermediate course must be followed: just as one must not rely too much on abstract knowledge when dealing with a new case, one also can not rely too much on intact case-based reasoning, when the latter is taken to mean reasoning to a new case from a single precedent case. Accordingly, as we have already mentioned, Cognitive Flexibility Hypertexts use segments drawn from cases--mini-cases--as the focus of instructional organization. KANE uses very short pieces of larger scenes as the mini-cases. Each mini-case receives interpretations across all of the multiple thematic dimensions that are relevant to it. (It is important to note that the entire case is always reconstructible in the hypertexts, and at some point during instruction the student should combine the mini-cases that belong to a larger case, so that they can learn how mini-cases configure to make up an intact case. The programs make it very easy to do this.)

It is worth calling special attention to a serious potential misinterpretation of what we mean by a mini-case. Mini-cases are not cases decomposed into their constituent features--they are not "abstract slices" of a case (e.g., the parts of a medical case that relate to vascular impedance). Such an approach would be antithetical to Cognitive Flexibility Theory since it would convey the mistaken notion, which is eagerly accepted by students, that the features of cases are independent, that one can study the aspects separately and then additively reassemble the whole case from those separately considered conceptual parts. Rather, mini-cases are chronological segments of a case (e.g., the first three hours of a battle, or frames 254-416 in a film). Thus, they retain some of the complex multiplicity found in the case as a whole. In a sense, they are microcosms, cases in miniature rather than separate case "compartments."

Advantages of mini-cases. There are several advantages to using mini-cases as the primary instructional organizing unit. We will discuss them in turn in order to show how design decisions in our hypertexts are bolstered by converging theoretical considerations:

1. Experience-consolidation and the mini-case: Accelerated acquisition of experience by compacting and elaborating cases. In an ill-structured domain you need to see lots of cases--the more case experience you have, the better your performance will be (especially independent transfer performance). For example, that's why experience is so valued in the professions (such as medicine); professional domains are notoriously complex and ill-structured, and professional training can not possibly provide sufficient experience in the limited amount of time available. Even with "problem-based" curricula, only a small number of cases actually get covered (which is very helpful for teaching the processes of intact case analysis, but less helpful for providing the multidimensionally analyzed, criss-crossed knowledge structures that need to be derived from case experience).

By breaking full cases into several mini-cases, and then conducting a rich analysis of each mini-case (i.e., compacting and elaborating case experience), and by focusing mainly on knowledge structural characteristics of each case, many more case experiences are provided, in an instructionally reasonable amount of time. The process of acquiring experience is consolidated. Again, better performance in ill-structured domains requires more cases and more case-processing experience.

It might be argued that the advantage of covering much larger numbers of cases by using mini-cases is lessened by the quality of the coverage. In fact, the learning made possible by our

richly analyzed mini-cases is quite potent. Consider all of the lessons taught in the context of just the one 45-second scene and few minutes of processing of commentary described earlier (the "gas in the lamp" scene). In that scene, it was clearly shown how four themes were simultaneously relevant, important lessons about the general nature of each of those themes were taught, illustrations of how the general themes had to be tailored to that particular scene (case of application) were presented, connections were drawn between that use of the themes and others in the film, and so on. All in less than five minutes, and all in the context of a real application of the conceptual/thematic knowledge that is being taught (not an artificial, contrived textbook example tailored to teaching the concepts but not representative of the types of cases that the concepts will eventually have to be used to understand--here, real scenes from the film). Now imagine what would happen if the in-depth understanding derived from that one mini-case presentation were to be compounded by similar rich lessons for ten other mini-cases, with each one related to the others. And that would take approximately one hour. The effect of dozens of these experiences over several days is obviously quite potent.

Again, however, it is worth pointing out that we do not claim that learning with our mini-cases substitutes for actual case experience. Rather, we contend that one important thing that comes from case experience can be more effectively conveyed by covering compacted and elaborated mini-cases in the way that we do: namely, the criss-crossed, multidimensional representation of the structure of case-based knowledge. Furthermore, these structural characteristics are taught more efficiently in our systems than in real-world case experience for a number of different reasons: conceptual structure is highlighted for the case, rather than having to be inferred; optional expert guidance is available; one is not dependent on serendipitous occurrences of instructionally useful cases in fortuitous sequences; etc. To get a lot of useful case experience, you don't have to actually physically experience the case. Again, if the fullness of the case's structural characteristics is successfully extracted, the effects may be better than actually "being there", because of the larger number of cases that can be covered, because of sequencing for instructional impact, and so on.

Summarizing, an important function of Cognitive Flexibility Hypertexts is that they consolidate the process of experience acquisition. By using mini-cases, the student sees many more examples of rich case analysis (e.g., of how concepts interact in a single case), in a much shorter amount of time. Each mini-case scene in KANE is a case-based demonstration of the processes of complex thematic analysis. Furthermore, each scene provides a rich lesson in the complex thematic structure of the film. And it deals with the conceptual information in the way that it is needed in ill-structured domains: namely, concepts are embedded in "practice"--the treatment of concepts is tied to the cases they are being applied to (and remember, in an ill-structured domain concept application is far more variable across instances, and thus more difficult, than in well-structured domains).

2. The mini-case and the problem of early introductions of complexity: Making complexity cognitively tractable for the learner. The use of mini-cases not only allows more cases to be covered, it also has beneficial effects on the cognitive manageability of the complex case instruction required in ill-structured domains. We have referred to our research findings that indicate difficulties with the traditional instructional approach of incremental additions of complexity (i.e., "start simple and then get more complex"). The early simplifications create impediments to the mastery of complexity introduced later (Feltovich et al., 1989; Spiro et al;

1989). This would suggest that complexity needs to be introduced earlier. The problem, obviously, is that complexity brings with it learning difficulties--you do not want to overwhelm or confuse the learner. Hence, a great advantage of the mini-case: it allows the earlier introduction of manageable complexity by presenting it in "bite-size" chunks.

In other words, Cognitive Flexibility Theory leads to a reconceptualization of instructional incrementalism. Instead of going from simple to complex treatments (with the attendant problem of creating barriers to advanced learning for the specific topics initially oversimplified, and the more general problem of instilling in students the mistaken belief that knowledge is organized more simply than it actually needs to be), instruction starts with complex treatments but situates them in cognitively manageable mini-cases. The extra cognitive difficulty of having to process multiple thematic representations of the same case is lessened by only having to understand and apply the multiple representations to a small case. So the student learns from the outset that the cases they will have to apply their knowledge to are complex (in that they require that multiple aspects of their knowledge representations be simultaneously and interactively superimposed) and they receive an easily graspable set of lessons about how some specific conceptual themes get instantiated in a particular context. A complex lesson in the application of conceptual knowledge is made much easier if the site for the application is as easy to grasp as, say, a 45-second scene from a film.

So, instead of making early instruction on a topic manageable for learners by oversimplifying the content (which will be counter-productive for more advanced knowledge acquisition), hypertexts like KANE present the conceptual complexity necessary for applying the knowledge to actual cases; but they do it in bite-size chunks that will help the students to avoid being overwhelmed.

3. Avoidance of maladaptive over-reliance on prototype cases. One thing that makes a domain ill-structured is that the processing of some new case is unlikely to benefit from a direct mapping (of structure or content) to any single prior case serving as a prototype. The individual case is complex, and there is considerable variability across cases. This is not to say that there is nothing to be learned from earlier cases. Far from it. Cognitive Flexibility Theory is a case-centered (rather than knowledge-centered), and cases are not just important to learn from: they also have a crucial role as precedents in the processing of new cases. But there is likely to be more than one antecedent case that will be useful for processing a new case. The Cuban Missile Crisis, as it was developing, was partly like the Appeasement of Munich, but partly unlike it, partly like Korea and partly unlike it, and so on (despite the fact that the policy arguments on President Kennedy's crisis team tended to revolve around which of the similar antecedents should be the precedent used in forming strategy for the United States' response). By using mini-cases, the monolithic integrity of the intact case is undermined, with the result that the tendency to over-rely on an exact mapping to just one precedent case is considerably lessened. Correspondingly, the need of mapping to multiple precedent/prototype sets, drawing on those parts of prior cases that are relevant is made clear, and the ability to appropriately assemble pieces of several precedent cases is enhanced.

4. Principled fragmentation for adaptive flexibility: The importance of mobile, recombinable knowledge elements. A key feature of ill-structured domains is that a single prepackaged schema or prototype case will typically be inadequate as background knowledge to

support the processing of a new case. Thus, intact schema retrieval (or prototype retrieval) as a knowledge-based processing mechanism must be replaced by situation-specific schema- and precedent-case-assembly. In a complex and ill-structured domains, small bits of information (either about prior cases or about abstract concepts) recombine in a large number of ways in the new cases that prior knowledge must be applied to. The use of mini-cases allows for easier situation-dependent knowledge assembly, because the cognitive processes of knowledge compilation are much easier to execute if they are operating on, say, 6 mini-cases rather than 6 full cases.

More importantly, however, availability of a large number of mini-cases permits a greater range of potential precedent-case assemblies for use in understanding new situations. Paradoxically, decomposition into very small discrete units enables a wider range of nondiscrete representations (similar to the paradoxical effect of breaking things into many very small units to approach a continuum in integral calculus). By using very small units that are then recombined with several others, it is much more likely that you will be able to assemble a precedent set to fit the many kinds of new situations that will be encountered in an ill-structured domain (a process facilitated by the mental record of and experience with recombination derived from past "criss-crossings"). A domain must be substantially deconstructed in order to have a wide range of possible reconstructions. Think of the way the use of small dots in comic strips allows more flexibility for developing different kinds of shadings than do bigger dots. In an ill-structured domain, cognitive flexibility to adapt to case variability requires a rich variety of potential shadings of the knowledge that is to be assembled. By structuring with small case segments you will be able to better match knowledge representations to the complexity of the world.

5. Retaining openness of interpretation. Meaning in ill-structured domains is multivalent. The goal must always be to resist the premature "closing down" of the interpretive process as soon as one account is identified. This is especially true in literary domains where the typical literary work will support many interpretations. Structuring over small segments (i.e., using mini-cases) helps to retain openness and plurality. Analyses based on large units will tend to stress elements found across the whole unit, and thereby miss elements that have more local occurrence; this will be true to a greater extent for domains that have complex cases (as in literary domains). Similarly, the larger the unit of analysis, the more that interpretation will tend to narrow to some "common denominator" that fits all the constituents, again missing the more irregularly occurring local complexities. Since cases tend not to be homogeneous in an ill-structured domain, representing at the mini-case level helps to avoid these reductive biases and thus offers greater opportunities for representational fidelity to a heterogeneous reality.

Barthes has made similar arguments. In his analyses of literary works (e.g., S/Z, 1974), he made the short segment, which he called a lexia (ranging in length from parts of a sentence to several sentences) the organizing unit for his analyses. He argued persuasively that this would result in an "open" and "plural" text, whereas if the text were structured over larger segments, it would be more likely to become closed to alternative interpretations. This is because many of the bases for opening up interpretation would be found only in parts of the text, but would be lost when the "grain size" is the text as a whole. As a result, the text would tend to have a more singular interpretation, uniformly applied throughout (again, because much

of the basis for realizing that the text has to be seen in multiple ways would have been narrowed away).

The avoidance of closed and narrow representations by structuring in small case segments is particularly important for adaptive flexibility in schema assembly, which requires rich diversity in the way cases are represented. If narrowing occurs in case representations, then you will have fewer "shadings" available to you to optimize the fit of prior knowledge to new cases.

6. Power and efficiency of the program. Besides the cognitive and instructional advantages already discussed, the use of the mini-case also has several "logistical" benefits for the development and operation of the computer hypertext program. For one thing, by having a larger numbers of mini-cases, each coded with a ten-slot vector with three possible values in each slot, the number of case contrasts and traversal routes for "landscape criss-crossing" that can automatically be generated by the program (without pre-programming of instructional sequences or pre-storing links across mini-cases) grows exponentially.

A time-saving advantage in hypertext development is that fewer cases have to be produced when each case results in several mini-case units. As we have already seen, quite a bit of instructional mileage can be gotten out of even a few cases if they are segmented into many complexly analyzed mini-cases. Of course, the efficiency of the program is further increased by the fact that Cognitive Flexibility Theory requires that each mini-case be used more than once (i.e., the same landscape site, or scene in this case, is re-visited on different traversals of the landscape).

Also, by using mini-cases it is much easier to effect a connection between a small part of one case and a small part of another case to which the former case is instructively related. It is more difficult to accomplish this when monolithic case blocks are used as organizing units.

7. Mini-cases help to avoid two common problems of hypertext: a confusing labyrinth of connections and the need to pre-store links. These technical issues in hypertext methodology will be addressed at the end of the paper.

We have devoted considerable space to the motivation for just one design choice in Cognitive Flexibility Hypertexts (including KANE), and the advantages that accrue to that choice: the principled decision to use mini-cases (short scenes in KANE) as the primary organizing unit of instruction. In the sections that follow we talk about other design decisions; although these are also considered to be crucial, we do not illustrate their rationale and benefits as exhaustively.

4.4.6.2 Multiple Knowledge Representations and Theme Selection to Maximize Transfer

In an ill-structured domain no single schema will provide sufficiently complete coverage, will account for sufficient variability in the way things happen in the domain. With the more limited view that results from a single representational system, the learner is not prepared to apply his or her knowledge to those new situations that are less relevant to that representation--and in an ill-structured domain there will be many such situations by

definition. Or, in the less extreme situation where multiple schemas are provided, it usually happens (either through instructional influences or because of learner biases) that one or two of the schemas assume precedence over the others. Although this is not as serious a problem as that when only one schema is provided, it is nevertheless the case that in an ill-structured domain any reduction in the operative perspectives for analyzing the wide variety of cases encountered will be disadvantageous. To make up for the inadequacies of any single representational perspective, additional perspectives must be added. With a sufficient number of perspectives, a fuller, more "three-dimensional" view of the domain is achieved. By providing 10 co-equal, frequently applicable themes in KANE, we are making it clear to the student that any one theme permits only a limited view of the landscape. The student learns that the abstract conceptual world for the ill-structured domain is complex and not easily simplified.

So, advanced cognition in complex and ill-structured domains requires multiple knowledge representations. As we have already indicated, each of the ten themes used in the Kane program is more than a mere descriptive attribute--each is a schema in itself. Each could be argued to provide the best full account of the Kane character. Rather than trying to show which is best, the program illustrates how they are all correct, and how their joint consideration produces the most adequate account. You can not impose one interpretive scheme on the film; any scene mixes and blends several interpretive schemes. This lesson is amply taught in KANE, as each mini-case (scene) is viewed from several thematic and symbolic perspectives. By helping the student to fully cover each case/scene by pluralistically representing it, transfer is fostered in several ways: (1) the student learns how to fully interpret cases, facilitating the full interpretation of new cases encountered later; (2) the multiple coding of cases provides a larger number of access routes for their later retrieval from memory as background knowledge precedents for understanding new cases; (3) the interaction of conceptual perspectives is taught by their simultaneous consideration within a single mini-case context; and (4) having 10 wide-scope themes will "cover the landscape"--it gives you more flexibility in tailoring for schema assembly. Also, having 10 themes helps students to avoid the reductive bias of a "uniformity of explanation" (Coulson et al., 1989; Fellovich, et al., 1989): with many equal themes none can dominate the others, and therefore all of these (now more minor) perspectives can be readily available to make their contribution as befits some new case.

4.4.6.3 Re-reading and rearrangement: Repetition without replication

The Kane program, like the other Cognitive Flexibility Hypertexts, relies heavily on the repetition of case information. However, this use of repetition is completely opposed to the typical purpose and consequences of repeating information in learning and instruction. Conventional uses of repetition in instruction are intended to strengthen the learning of some aspect of knowledge, with successive presentations each intended to mean the same thing each time. In contrast, our use of repetition is non-replicative. The aim is to illustrate the complexity of case information. Following the metaphor of the criss-crossed landscape, the same content is presented in different contexts. This helps to keep interpretation from rigidifying towards an overly narrow subset of the lessons that the content should be teaching (and that the learner must be prepared to utilize in transfer situations). That is, the repeated presentations aim to point out for students how the same case information can take on

importantly different shades of meaning at different times and how each case has many facets, some of which will tend not to be noticed in any single context of occurrence. In complex and ill-structured domains, each unit of content is multifaceted. Each presentation of that content in a different context highlights another aspect of that multifacetedness, as well as illustrating context-dependencies. There is a limit to the number of lessons that can be learned in any single presentation of content material, in any single context of presentation. Furthermore, by presenting the same case information at different times, in the context of various other cases, and with different conceptual elements stressed, a web of case and context interrelationships of the kind necessary for flexible knowledge assembly and transfer in ill-structured domains is established.

Relationship to poststructuralism in literary theory and philosophy. A similar role for re-reading in order to avoid reducing complexity was a central feature of Barthes' (1967) analysis of Balzac's short novel, "Sarassine," in *S/Z*. In fact, our general approach has many affinities to poststructuralist literary theories (like that of Barthes in the late 1960s), which also stress such factors as multiple codes, the importance of knowledge fragments, and the nonunifiability of rich cases by any single unifying logic. Not coincidentally, Wittgenstein's later work, which strongly influenced our thinking, also has been adopted by many poststructuralists. (Interestingly, Wittgenstein had originally intended to have an elaborate set of cross-references placed below each sketch in the Philosophical Investigations to enable criss-crossing the landscape of sketches in various directions with frequent re-readings in new contexts.)

4.4.6.4 Theme-Based Exploration: Teaching Conceptual Complexity and Variability in Conceptual Application

One of the main features built into KANE is the ability to have the program search for occurrences of any chosen theme and then re-edit the film to show just those scenes found in the search (along with accompanying commentary, stage-setting information, etc.). Theme-based explorations of the film are intended to teach both the complexity of the themes and the nature of their variability of instantiation in actual scene contexts. An important feature of abstract concepts in ill-structured domains is their irregularity of application (see also Barsalou, 1987). That is, each concept is used in a lot of different ways. Also, the same concept will apply to a variety of kinds of cases. These features have the consequence of making it hard to: (1) go from cases to concepts (problems of knowing which concepts are relevant from the case information), and (2) go from concepts to cases (problems of tailoring the way a concept is used to a particular context of its application). [Once again, in this chapter we use the term concept very broadly to make a gross distinction between abstract knowledge and case information. With this loose criterion, a complex theme in KANE can be referred to as a concept.]

Theme-based search and variability of instantiation. In general, the concepts of an ill-structured domain can not be transmitted to learners the way they are in well-structured domains, namely by some direct process of providing general principles or definitions (perhaps with one or two examples as illustrations of the general principles). Instead, knowledge of concepts comes from having their uses (instantiations) demonstrated, rather than the concept being specified in the abstract. In an ill-structured domain the meaning of concepts is implicit,

at least partly, in the pattern of its uses of the concept (Wittgenstein, 1953). The theme-search feature of KANE (and a similar feature in the other Cognitive Flexibility Hypertexts) does what must be done in an ill-structured domain: it allows the nature of concepts to be shown (by a guided presentation of their actual occurrences), not told. The nature of the variability of theme-instantiation in KANE is conveyed in part by sequentially viewing an entire series of scenes that are instances of the theme. For example, the user can ask to view a series of brief scenes that all illustrate the Hollow Man theme. Thus, the program shows the variability in conceptual use--you see a set of different cases that are said to be instances of the concept. (Of course, you also see what is in common across the uses--but this is much easier to teach.)

Commentary and expert guidance: Tailoring, integrating, and preparing for access. The thematic commentary that accompanies scenes also helps to convey conceptual complexity. Commentaries provide several kinds of assistance. First, they explain how the scene instantiates the theme; that is, they tailor the theme to the case. Second, they discuss features of the scene that lead one to consider it to be an instance of the theme, information that helps students learn the difficult process of accessing concepts from case information. Third, the commentaries help with two kinds of integration functions: they relate the theme to other themes in the same scene (building theme connections and reminding learners about the importance of the simultaneous use of multiple perspectives); and they relate the current instance of the theme to other instances of the theme elsewhere in the film.

Also helping in the management of conceptual complexity is the use of more differentiated sub-theme designations that tell what type of instantiation of the theme is being observed. That is, each theme is analyzed into various "senses," and the instantiated sense is presented in the header for the commentary. For example, a sub-theme of the Power theme is "Kane controls others by having them depend on him." The sub-theme information also accompanies the viewing of a scene (i.e., whenever a theme is instantiated in the film, that theme's name is presented in colored print that overlays the film, accompanied in parentheses by a sub-theme designation). Of course, the problems of ill-structuredness apply at the level of sub-themes as well: sub-themes lack definability in the same way that concepts do and, like ill-structured concepts, must have their uses shown or demonstrated.

A note on "situated cognition" and the relative importance of conceptual knowledge. A popular call has been raised recently for "situated learning" (e.g., Collins, Brown, & Newman, in press). The emphasis that has been placed on situations at the expense of conceptual knowledge has led some people (although not Collins et al.) to suspect that abstract conceptual knowledge is relatively unimportant, that all that matters is the situations or cases in which learning occurs. The emphasis in Cognitive Flexibility Theory is clear: concepts and cases are both essential. However, conceptual knowledge must be taught in the context of actual cases of its application (not "in the abstract"), and the ill-structured nature of the use of conceptual knowledge must be acknowledged and directly addressed in theories of learning, knowledge representation, and instruction. The approach of Cognitive Flexibility Hypertexts is intended to effect an integration of conceptual and situational learning, in which each is appropriately thought about in terms of the other.

4.4.6.5 Theme-combination-based Exploration: Avoiding Compartmentalization and Assembling Higher-Order Knowledge Structures

An important feature of Cognitive Flexibility Hypertexts is their ability to search for mini-cases that are instances of some combination of abstract concepts. For example, the KANE program has a menu option that allows for traversals of the film (i.e., re-edited scene juxtapositions) to be selected that highlight combinations of themes.

Non-insular treatment of themes. A common problem in traditional instruction is compartmentalization: conceptual areas that are highly related are presented in separate chapters, lessons, classes, and so on. As a result, the knowledge ends up being represented as if it were in separate compartments. When knowledge from across compartments later has to be combined for use in some situation, the representational basis for the conceptual combination is weak. In ill-structured domains, conceptual combination is the rule rather than the exception. Thus, by allowing the film to be explored as a function of theme-combinations, students learn about the patterns of interaction of conceptual themes, their context-dependencies (i.e., the way theme meanings are altered by the context of other themes that they occur with), and so on. These lessons (as well as KANE's approach to teaching several themes in the context of each mini-case scene) help to vitiate the force of the compartmentalization bias and to provide knowledge and skill for processing conceptual information in a non-compartmentalized way.

Thus, not only is there a multiplicity of themes, but they function as something more than a list of independent items. As befits an ill-structured domain, the themes are not treated as insular, separable "compartments"; rather, the themes are shown to interpenetrate, to have complex patterns of mutual dependence upon each other.

Schema-building. Another use of the theme-combination feature is in generating hypotheses about complex structural models of the film, and then enabling explorations to test and refine those models. These hypotheses can be triggered by such things as chance observation of patterns in the haphazard exploration of theme co-occurrences (i.e., noticing two themes going together in a few different scenes). Or they can be explicitly suggested by the film itself (as in the example presented in the second paragraph below). Or, by exploring small combinations of themes, larger explanatory structures are suggested (also illustrated below). These hypothesized larger structures can then be explored using theme-combination search to validate the hypotheses. This kind of activity provides practice in situation-sensitive schema assembly (e.g., preparing to write an essay, as in the example in the second paragraph following), as well as helping to build "mini-schemas"--complex theme combinations that have partial applicability.

An example of a mini-schema would be the combination of the Wealth Corrupts and the Outsized Ambition themes, as in "Kane's tendency as he got older to buy things instead of earning accomplishments (Wealth Corrupts theme) interfered with the attainment of his earlier ambitions (Ambition theme)." One could then test this mini-schema or examine its variability across instantiations in the same way that variability for a single theme is learned, namely, by juxtaposing a set of scenes that constitute instances of the theme combination.

A lengthy example of using KANE to build a higher-order schema (a complex model) is the following sequence of steps using the program to answer an essay question: Why did Kane fail to achieve the ambitions which seemed within his grasp during his more youthful years?

A start towards assembling knowledge to answer this question might come from the mini-case scene that immediately follows the one presented in detail earlier in the chapter (the "gas in the lamp" scene). After Kane says that he wants his newspaper (which he equates with himself) to be "as important to New York as the gas in this light", his friend Leland asks him how he is going to achieve that goal. Kane goes on to say that he will do it by putting a "Declaration of Principles" on the front page of the paper. He will be honest with the citizens of New York, he will look after the interests of the disenfranchised, and so on. His associate Bernstein pointedly responds by telling Kane "You don't want to make any promises you're not going to be able to keep. Kane says that he will keep these promises. That is, he will achieve his ambition for the newspaper (and thus himself) by principled behavior. This suggests a possible relationship between the Ambition theme and the Principled Behavior theme. Another part of the puzzle might suggest itself if the student recalls another mini-case scene: at a difficult moment Kane states with some gravity that he could have been a great man if he had not been very rich. Summarizing thus far, we have three themes that might figure in an answer to the question about Kane's failed ambitions: the themes of Principled Behavior, Wealth Corrupts, and, of course, Ambition. The hypothetical model for responding to the essay question at this point might have the form "Kane's ambition was thwarted by his wealth interfering with the path he had chosen to attain his ambitions, principled behavior."

This suggests a re-editing of the film to look at the joint occurrences of these three themes. After the program executes the search, among the scenes encountered will be one mini-case that is especially provocative. The newspaper has become a huge success, and Kane and the staff are celebrating at an extravagant (and mildly decadent) party. Kane is about to embark on a vacation trip to Europe, and Bernstein stops the party to ask Kane if he intends to buy a lot of statues while abroad. Kane says yes, and then Bernstein keeps saying, with emphasis as indicated, "Promise me, Mr. Kane. Promise me." There is no apparent reason for Bernstein to go on and on about Kane promising him that he will buy those statues (at least when viewing the film sequentially). However, in the context of having recently encountered the scene in which Bernstein pointedly suggested to Kane that he not make an promises he was not going to be able to keep, a clear link is established to the principled behavior that was mentioned in that scene as what was going to make the newspaper and Kane great. Once that allusion to the earlier scene is established, it becomes especially significant that Kane responds to Bernstein: "Yes, I promise you. But you don't expect me to keep any of those promises, do you?" Clearly, the whole unusual interchange about promising (as well as Bernstein's earlier unusual emphasis on Kane not making promises he could not keep) is intended to signal a connection between the two scenes (or at least supports such a connection, if it was not intended), a connection that is portentous for the theme of Kane's failed ambition. What might have seemed to be a trivial (though odd) statement about making promises now takes on an enriched meaning by its juxtaposition with a

more serious scene that also had an odd emphasis on promising (tied to Kane's promise of principled behavior as a way to make him and his newspaper great). Furthermore, the point about Kane not keeping any of his promises is made in a context in which Kane has begun to clearly show that he is allowing his wealth to corrupt him (i.e., in several scenes, obvious connections have been drawn which tie unprincipled behavior by Kane to the corrupting influences of his wealth, as revealed by having the computer do a two-theme search. For example, he repeatedly indicates that he thinks he can act in an unprincipled manner towards friends and then buy them off.

Indeed, after the "promise me" scene at the party, Kane does begin to slide downhill, his lofty ambitions fading further and further from possibility with the passage of time.

Thus, the KANE program provides a substantial start toward a cohesive interpretation of Kane's failed ambitions in terms of his renunciation of principled behavior because of the corruption of his wealth. Especially striking is the way the program helps to make connections that would otherwise be very difficult to make. Of course, this is only one of many possible responses to the essay question. What is important is that the KANE program helps students to respond to situational demands in a manner that is original, complex and warranted (i.e., the program helps the student to assemble justification for a budding complex interpretation). After a short time working with the KANE program, very intricate weavings of themes, of the kind just illustrated, become very straightforward.

4.4.4.6 Sequencing to produce cognitive structures with woven interconnectedness

A topic that we will not treat in detail in this chapter and that is not yet fully implemented in the Cognitive Flexibility Hypertext systems is sequence and arrangement. Following Wittgenstein (1953), we consider the strength of representations in "structured domains to depend not on a single thread running throughout--that is too reductive of the domain's complexity--but rather on the overlapping of many strands. This criterion increases the possible hinges for schema or precedent assembly, the number of access/retrieval routes in memory, and so on. So, it is not just important that a domain's landscape be criss-crossed--how it is criss-crossed is an important consideration. Of course, it should be criss-crossed in such a way that useful instructional contrasts are highlighted. Furthermore, sequential arrangements should be such as to promote knowledge representations characterized by highly woven interconnectedness along several conceptual dimensions, rather than highly compartmentalized structures with connectedness of many elements determined by some single organizational system.

One important sequencing principle that we will incorporate into our systems is that of "intermediateness": mini-cases should be sequenced for presentation in such a manner that two representational extremes will be avoided. First, a case should not be the next one presented if it would too closely parallel the thematic interactions in several recently presented cases; this would help students to avoid overgeneralizations, a major hazard in ill-structured domains. Second, at the other extreme, a case should not be so dissimilar from those that were recently presented that it would give the learner the mistaken notion that there is no role for conceptual abstractions that extend across cases; this would help students to avoid the perception that each

cases is unique. Rather than either of these extremes, cases should be selected because their strength of relation to preceding cases is intermediate, partially overlapping and partially nonoverlapping (see Spiro et al., 1987). This would promote an appreciation of the importance of both knowledge-based and case-based analysis, without either assuming precedence over the other.

The KANE program allows for both experimenter-controlled and student-controlled sequencing. The program keeps track of student-selected traversals. The relative efficacy of different sequencing patterns in the nonlinear presentation of materials, or of different patterns of spontaneous student exploration, is a central concern of our theory, and one that is deserving of considerable empirical investigation generally.

4.4.4.7 Other Features

There are many other features of the KANE program and Cognitive Flexibility Hypertexts that cannot be discussed in this chapter because of reasons of space (e.g., options related to user-customization; user- and teacher-construction of interpretive essays--for example, one has been built for the famous reference to Rosebud in KANE; and active learner participation). We have focused in this chapter on those features that cast the most direct light on the most important theoretical underpinnings of the instructional approach.

To take just one example of a potent feature that can not be discussed in depth, consider the "special initiative" option, which allows for the easy development and use of mini-case sequences that illustrate new themes not covered in the provided set. Obviously, this option makes the computer program more flexible. However, it also has some cognitive significance. For example, this option increases the program's power to deal with thematic combinations. Since the "special initiative" option presents specially programmed sequences that are accompanied by their thematic overlays, one can be thinking analytically about the various themes represented in a scene at the same time that a symbolic feature, the topic of the special initiative, is affectively "coloring" that thematic analysis. For example if snow is supposed to symbolically evoke feelings of contentment and security associated with childhood, that mood can be experienced in the background while one is consciously thinking about the role of the Wealth Corrupts theme in the same scene (just as, when you are angry, you can be thinking about something unrelated to your anger, but be doing it angrily--the emotion effortlessly overlays the other cognition). Thoughts about one topic can be colored by the feelings evoked by another topic, thus expanding the scope of possible literary understanding and mental representation (Spiro, Crismore, & Turner, 1982).

4.4.4.8 Generality of the Approach Demonstrated by Parallels in Cognitive Flexibility Hypertexts for Random Access Instruction Across Domains

We have stressed the point that the approach taken in the KANE program is based upon principles of random access instruction derived from Cognitive Flexibility Theory (Spiro et al., 1987, 1988). It is very important to make clear that that theory has the same implications for random access instruction in other, very different domains. That is, the systematic, theory-based approach in Cognitive Flexibility Hypertexts is domain-independent--essentially the same hypertext approach may be employed across ill-structured domains. There are currently

two other functioning hypertext prototypes besides KANE, one in cardiovascular medicine and one in military strategy. All three of these Cognitive Flexibility Hypertexts have the same primary features discussed in this chapter (as well as many others we did not have time to present here). All of them: are organized around mini-cases (e.g., parts of battles in the military strategy program); impose multiple conceptual perspectives on each mini-case; allow for concept-based search across cases (highlighting conceptual variability of application and instances of conceptual combinations); provide commentary on the relationship between conceptual and case elements (i.e., they teach the concepts in the context of cases and show how the concepts have to be tailored to the individual case); and so on. (The Cardioworld Explorer hypertext is discussed in Spiro et al., 1988. A detailed treatment of all of the hypertexts will be found in Spiro et al., in preparation, a paper about hypertext generally and about our approach specifically.)

However, by claiming domain-independence for our hypertext approach we are not saying that we dispute the widely reported findings of domain-dependence in cognitive science. In fact, there are extensive differences in the nature of knowledge and in the way the knowledge must be used between the domains of literary interpretation, biomedical cognition, and military strategy. However, as we mentioned before, it is not the same thing to say that two landscapes are radically different as it is to say that there are important principles of landscape exploration that are general across domains (e.g., multiple representations; repeated presentations in different contexts; etc.)--it is principles of the latter kind that form the basis for the commonality of Cognitive Flexibility Hypertexts across domains. When the specific content of a domain is used to instantiate the general frame that underlies Cognitive Flexibility Hypertexts, domain-specific features emerge. Again, the fact that the same primary features were successfully incorporated into hypertexts for three highly dissimilar domains attests to the generality of this theory-based instructional approach.

4.4.5 Benefits of the Cognitive Flexibility Approach to Random Access Instruction: Developing Knowledge Representations and Cognitive Processes Appropriate to Advanced Learning in Complex and Ill-Structured Domains

The application of Cognitive Flexibility Theory to advanced instruction using computer hypertexts, as illustrated in this chapter by the KANE program, results in several major benefits. These can be briefly summarized as follows:

4.4.5.1 Towards a New Incrementalism: Tractability of Learning and Instruction Without Inhibiting the Acquisition of Complexity by Early Oversimplification

The attention to covering necessary complexity is not achieved at the expense of overwhelming and confusing the learner. By using bite-size chunks of complexity in the early stages of advanced instruction (mini-cases), followed by the use of more complete cases later, instruction that addresses complexity may be introduced from the outset because it is staged in the context of a limited and manageable (but nevertheless moderately rich and ecologically representative) example--i.e., the use of mini-cases reduces the cognitive demands on the learner without the hazards of oversimplified "concepts-first" approaches. (A number of other

features of our approach also help learners to manage the complexity that is presented, including the various options available to the learner to get theme commentary and guidance, stage-setting information, scene re-viewings, verbal overlays of active themes during film viewing, and so on.)

4.4.5.2 Advanced Learning in Instructionally Reasonable Amounts of Time: Accelerated Acquisition of Case Experience

Another problem associated with the need to cover complexity and to cover a large number of cases (a prerequisite of advanced performance in ill-structured domains) is the amount of time required. It would be desirable to achieve these goals in the time available for instruction (i.e., a course, or a program of study). The process of acquisition of complex case experience is accelerated in our approach by the use of elaborated mini-cases and by making use of the same, increasingly familiar mini-case in different contexts (thus eliminating the need to spend time learning as many new cases). Thus case experience can accumulate far more rapidly with this approach than it does in either: (a) case-based curricula that spend much more time on each case (again, you need to see lots of cases in ill-structured domains because of the many different forms cases assume), or (b) the natural exposure to case experience (which is haphazard, and thus tailored to instructional needs, nor guided as to properties of conceptual structure).

4.4.5.3 Teaching Concepts-in-Practice: Avoidance of Over-reliance on Knowledge-Based or Case-Based Representational Extremes

The way that advanced topics are usually taught in instructionally reasonable amounts of time is by stressing abstract conceptual knowledge at the expense of exposure to cases. The hope is that the abstractions will have wide scope of application to new cases. In an ill-structured domain, this is a vain hope. For one thing, in an ill-structured domain, concepts vary too much in the way they apply to cases. Cognitive Flexibility Hypertexts like KANE directly address conceptual variability across cases of conceptual application. And these hypertexts situate the teaching of concepts in the context of actual cases, demonstrating how concepts are tailored to cases. Thus, our programs neither neglect cases to teach concepts, nor concepts to teach cases--both are taught in the context of each other. Learning is situated, but abstract knowledge is not ignored. Our approach teaches concepts and cases simultaneously, not separately: concepts-in-practice.

4.4.5.4 Avoidance of Counter-productive Compartmentalization

A serious problem in the preparation for knowledge transfer by traditional instruction is the presentation of information in highly compartmentalized forms. Different conceptual topics are treated in different parts of texts and at different times. When cases are presented, they are usually dealt with in isolation and very seldom related to other cases. But in an ill-structured domain, knowledge can rarely be used intact--parts of topics/schemas/concepts must be combined to form schema-assemblies, and cases must be combined with other cases to form precedent sets. These processes are inhibited by the compartmentalization of knowledge representation that results from compartmentalization in instruction. In our approach, compartmentalization is avoided in a number of different ways. Several themes are

simultaneously considered when each mini-case is being processed. Search for examples of any possible theme combination is facilitated. Commentary on one theme will contain allusions (cross-references) to other themes.

Similarly, the likelihood of compartmentalization by cases is minimized. An intact case is less likely to be rigidly treated as a monolithic entity after having been broken into several mini-case segments. Also, the multiple conceptual codings of each mini-case causes it to fall into several conceptual categories. And the ubiquity of case juxtapositions when using the KANE program, as well as the allusions to other cases in the theme commentaries and the thematic overlap across cases, are all designed to build interconnections among cases, not separation.

4.4.5.5 Teaching Situation-Adaptive Knowledge Compilation: From Intact Schema Retrieval to Schema Assembly

Schema and precedent assembly, as opposed to intact schema retrieval, is, again, a crucial tenet of Cognitive Flexibility Theory. You can not have a prepackaged schema for every situation you will encounter in an ill-structured domain. A new case will be kind of like an aspect of one prior case, kind of like an aspect of another prior case, and so on. Similarly, for the relevance of parts of different concepts, appropriately assembled to fit the new case. So you need to build schemas to fit new situations. This requires flexibility in knowledge representations. For flexibility you need many movable/recombinable knowledge elements (meaningful fragments of knowledge--"partial theories") within a web-like structure. This is accomplished in Cognitive Flexibility Hypertexts by such features as the use of minicases/scenes (for later precedent assembly) and by having a fairly large number of wide-scope Themes/Schemas/Perspectives for conceptual schema assembly. The larger the available set precedent cases and conceptual perspectives, the greater the likelihood that you will have an optimal combination for dealing with the odd new case. The more pieces you have to work with that each make some nonoverlapping contribution compared to the others, the greater the adaptive flexibility you will have to respond to complex and changing case realities. As we said earlier in the chapter, an ill-structured domain must be substantially deconstructed (while retaining complex interactions of the parts) in order to have a wide range of possible reconstructions. And facilitating the adaptive assembly of these pieces of cases and fragmentary "theories" is the earlier record of and experience with criss-crossing them during instruction (e.g., assembly routes are suggested--you can get between more places in the landscape in more different ways by having a rich network of interconnecting routes).

So, for schema assembly you need: 1) lots of little pieces of reality (mini-cases) and of conceptual knowledge (multiple themes/concepts/partial theories), and 2) a way of assembling them, of putting them together to fit a new case. These needs are supported by features of the approach such as theme-combinations, repetition in new contexts, the more rapid accumulation of experience due to the use of compacted and elaborated mini-cases, and single theme search. Consider single theme search: it demonstrates that even routes with the same theme-name can be differentiated, which allows you to even better tailor your knowledge to highly diverse and complex new cases: you can pick just the right flavor of thematic connection--your palette for painting case reality has a much richer and more subtle range of colors. All of the other features similarly contribute to promoting schema assembly ability.

4.4.6 Addressing Some Problems in Hypertext Methodology

The work on KANE and the other Cognitive Flexibility Hypertexts contribute many new ideas to the growing literature on hypertext methodology. Two especially important ones will be addressed here. First, the use of hypertext is made less daunting for the student. Second, it uses a procedure that automatically generates connections, rather than having to have all links stored in advance. These will be discussed in turn. (Of course, we consider the most important contribution to be the theoretical one that is the subject of this chapter.)

4.4.6.1 How to Avoid the Problem of "Getting Lost in Hyperspace"

A common problem with hypertexts is that the user soon gets lost in a labyrinth of connections, and loses track of the sense of his exploration, as well as his or her physical place in the hypertext (collectively referred to as "getting lost in hyperspace"). The use of a case-centered instructional scheme with mini-cases as instructional and programming foci solves this problem. You can never get lost because you are never more than one connection from the focus of instruction. In a sense, each mini-case, begins a complete and independent unit of instruction. All departures for commentary, guidance, context-setting, and so on, take you right back to the case-at-hand. Each mini-case starts a new lesson (constitutes a new experience).

4.4.6.2 Latent Rather than Programmed Links: Multithematic Coding of Cases

In conventional approaches to hypertext development (e.g., those using HyperCard), it is usually the case that any connections between knowledge nodes (between "note-cards") that will be available for the user to explore have to be anticipated and explicitly built into the program. This is a limitation of most hypertext programs that does not apply to Cognitive Flexibility Hypertexts. By coding each mini-case with a vector of relevant themes, the KANE program automatically generates instructive case sequencings (and many times more of them than would be possible if all links had to be pre-stored--the number of mini-case juxtapositions is so large a number because very many mini-cases are used, and there are many values in the search vector for each of them). So the program can be used for much longer without duplication of instruction and with many more lessons being taught.

We have seen how connections are automatically identified by the program for presentation to the user. But the mere juxtaposition of cases (even with appended commentary) does not guarantee that important connections between them that are not explicitly drawn out will be represented in the mind of the user. How, then, do the connections that are implicit in the program get formed in the mind (e.g., connections across mini-cases)? The same way they are in actual experience. After all, the computer program merely presents experiences (albeit, stripped down to their structurally-relevant features). However, the experiences in Cognitive Flexibility Hypertexts are different from actual experiences in several important ways that make it easier to induce interconnected knowledge representations from exposure to cases: (1) the cases are immediately juxtaposed (hours or days do not pass between non-routine cases); (2) the cases are thematically related (whereas there is no guarantee of instructive relatedness across naturally occurring adjacent cases); (3) the cases are stripped down to structurally significant features, making it easier to extract dimensions of structural

relatedness; (4) the cases are accompanied by expert commentary and guidance; and, finally; (5) because the cases are short, they are each easier to remember and more of them can be presented in a short amount of time, facilitating the recognition of relationships across cases.

4.4.7 Conclusion

The overall effect of the features of random access instruction that are derived from Cognitive Flexibility Theory and embodied in "Exploring Thematic Structure in Citizen Kane" is a program that allows the Kane character to be viewed from a very large number of valid perspectives. The result of overlaying more and more points of view on the same content material (while at the same time reducing initial demands on the learner resulting from this extra complexity by working with easily digestible mini-cases) is a kind of "stereographic" representation--the multidimensional fullness of the content is increasingly approximated with each additional perspective that is presented. Furthermore, the theme combination feature of the program permits an incremental buildup of a picture of the interrelations among the thematic perspectives. Instead of just having a set of independent conceptual perspectives that have to be additively assembled, the complex pattern of their intertwinement in the actual cases (scenes) can emerge. By re-presenting the same information in different contexts and from different perspectives, the complexity of that information is made more resistant to oversimplification. As a result, knowledge representation is made more multidimensional--and knowledge that will have to be used in many different ways has to be represented in many different ways, with the potential to form various combinations with other aspects of knowledge as required by new contexts of knowledge use.

The result of instruction of this sort is deeper understanding of complexity and nuance, understanding that provides learners with a basis for going beyond what was explicitly taught. In ill-structured domains, there is considerable variability in the way knowledge has to be used across the set of potential knowledge application situations. Correspondingly, there is a greater burden on learners to be able to independently apply their knowledge, rather than relying on prepackaged "prescriptions" for knowledge application provided by teachers and textbooks. Therefore, it is essential that learners be presented with a cognitively tractable picture of the landscape of varieties of knowledge use. And learners must be guided in the development of cognitive skills for effectively traversing those landscapes to independently and adaptively assemble knowledge to fit the new situations that that knowledge must be applied to. These are the aims of the hypertext instructional systems that implement Cognitive Flexibility Theory. In Cognitive Flexibility Hypertexts (like "Exploring Thematic Structure in Citizen Kane"), these aims are achieved in an instructional environment that reconciles agendas whose seeming incompatibility would be expected to impose extreme obstacles: instructional material is presented in a manner that does not sacrifice complexity, yet takes an instructionally reasonable amount of time to cover, and does not overwhelm the learner.

A final point: The instructional approach described in this chapter is difficult. But sometimes advanced knowledge acquisition has to be hard. There are data that indicate that difficult instruction tends to be neglected, at great cost to learning outcomes (Feltovich et al., 1989; Spiro et al., 1988). The trick is to make advanced learning as easy as possible without sacrificing the integrity of the material to be learned. That is what Cognitive Flexibility Hypertexts attempt to do.

4.4.8 Footnotes

The principal contributor to this chapter was Rand Spiro, with additional contributions by Jihn-Chang Jehng.

¹The three hypertext systems are fully functioning prototypes. However, because they are currently used only for research, they contain a limited number of cases and commentaries. They are not fully developed curricula, although preparing them for that function would simply be a matter of adding more cases and commentaries--nothing in the computer programs themselves would change.

²One further note may be helpful. The terms hypertext, hypermedia, and HyperCard are often confused. Hypertext has already been discussed. HyperCard is merely one kind of programming environment for hypertext development; it was not used for the Kane program. Hypermedia refers to nonlinear computer learning systems in any medium (including multiple media). In that sense, it is a more general term than hypertext. However, since only one of the nonlinear learning programs that have been based on Cognitive flexibility Theory employs multiple media (the Citizen Kane program), we feel uncomfortable referring to the set of systems as hypermedia. Yet the set should have a common name, given that a common theoretical basis is shared. So we refer to all of the systems as hypertexts. This designation is consistent with the more general usage of the term "text" in post-structuralist literary theory to refer to any object of rich interpretation, including pictures. So, the instructional program using Citizen Kane is a nonlinear 'text' that relies heavily on film segments, as well as written text.

³An example of a transfer measure for Citizen Kane is comprehension of critical commentary on the film. It has been amply demonstrated that text comprehension requires the mobilization of an appropriate background knowledge scaffolding. This requisite background knowledge changes across the many texts written about the film. Ability to assemble or comprehension-supporting background knowledge across a wide variety of texts (compared to comprehension scores of control subjects) would indicate the acquisition of one kind of flexible transfer ability.

⁴A note on terminology will be helpful before continuing. In this chapter we use the term concept very broadly. We sometimes call schemas concepts, and we frequently refer to the complex themes of Citizen Kane as concepts. Sometimes abstract perspectives are called concepts. No theoretical significance should be attached to this usage. It is merely a convenient way to make a very gross distinction between abstract, conceptual knowledge and information about concrete cases. Similarly for our use of the term case. We refer to the short scenes from Citizen Kane as mini-cases (or sometimes, in more general theoretical discussions, as cases). Sometimes we call examples or events cases. We are simply lumping together as cases anything that is an actual happening or a description of an actual happening, whether it is a scene from a film, a medical case, or an historical event. This is part of the same fundamental distinction we draw between conceptual knowledge and cases (things that actually happen).

⁵The program was conceived by the first author, who also conducted the content analysis of the film. The second author wrote the program code from the specifications provided to him. Michael Jacobson helped with some of the user-interface design.

⁶It is worth noting again that the Cognitive Flexibility Hypertext approach is a general one: the same principles lead to very similar hypertexts in radically different knowledge domains. The issue of domain-independence is relevant here. We believe that there are considerable differences across domains, but the principles of domain exploration are the same (just as New England and Southwestern landscapes differ while general principles of how to explore a landscape can be applied to both). For example, in any landscape or complex knowledge domain, multiple traversals will be important.

⁷It should be noted that KANE also allows well-structured aspects to be taught (for example, besides showing how themes differ across instantiations, the program teaches what distinguishes each theme (i.e., what theme instantiations have in common). However, given limitations of space, we will emphasize the more novel treatment of complexity and ill-structuredness.

Chapter 4.5 Cognitive Flexibility, Hypertext, Computer-Assisted Instruction, and the Promotion of Transfer for Complex Knowledge

4.5.1 Overview of Chapter

Recent applications of educationally oriented hypertext systems have attempted to structure knowledge-rich *learning environments*, often in conceptually demanding domains and for students who are at an advanced stage of learning. A study is reported that investigated the use of recent cognitive instructional theory to structure a hypertext system to provide instruction in a complex and ill-structured content area. The results of the study revealed that the more hypertext-like treatment promoted superior *transfer* to new situations, while a hybrid hypertext/computer-based drill treatment led to higher performance on the measures of *memory* for factual knowledge. The experimental findings are discussed from the perspective of the cognitive theoretical framework that underlaid the experimental hypertext system. Instructional prescriptions for hypertext learning systems, such as the use of multiple representations of knowledge and the demonstration of critical interrelationships between abstract and case-centered knowledge components in multiple contexts, were found in the study to help promote the mastery of complex knowledge and the ability to apply this knowledge to new situations. Some implications of these findings for the development of hypertext learning environments are also discussed.

4.5.2 Introduction

While early applications of hypertext technology often focused on providing flexible access to a wide variety of information types, newer educationally oriented hypertexts have attempted to structure knowledge-rich *learning environments* (e.g., Beeman et al., 1987; Beeman et al., 1988; Conklin, 1987; Crane and Mylonas, 1988; Dede, 1987, 1988; Jonassen, 1986, 1988; Spiro and Jehng, 1990). An important characteristic of these recently developed hypertext systems is that, unlike traditional computer-based instructional approaches which are commonly developed in instructional contexts that either are at an introductory stage of learning or are in more conceptually well-structured content areas, these hypertext learning environments are frequently developed for more advanced learning situations involving complex and cognitively challenging knowledge.

Although hypertext systems have been used recently to promote the achievement of more advanced learning goals, there has been little discussion in the hypertext literature concerning recently identified learning problems associated with more demanding instructional contexts (Jacobson and Spiro, 1991a; Spiro, Feltovich, Jacobson, and Coulson, in press; Spiro and Jehng, 1990). Serious learning failures have been identified in recent cognitive instructional research, such as the development of conceptual misunderstandings (e.g., Feltovich, Spiro, and Coulson, 1989; McCloskey, 1983; Perkins and Simmons, 1989; Spiro, Coulson, Feltovich, and Anderson, 1988) and the inability of students to transfer knowledge to relevant new situations (e.g., Bransford, Franks, Vye, and Sherwood, 1989; Brown, 1989; Gick and Holyoak, 1987; Schoenfeld, 1983). Unfortunately, many instructional applications of hypertext have been developed in an atheoretical manner and have often not considered learning problems such as these that confront learners of complex knowledge (Jacobson and Spiro, 1991a; Spiro et al., in press; Spiro and Jehng, 1990).

The major goal of this chapter is to discuss features of a theory-based hypertext learning environment employed in a study that investigated knowledge acquisition and transfer issues in a complex and ill-structured domain. The results of this study revealed significant knowledge transfer effects associated with a new approach to structuring an instructional hypertext system. As a detailed account of the research methodology and findings may be found in Jacobson and Spiro (1991b), this chapter will focus on the theoretical framework used to structure the hypertext system that contributed to the attainment of knowledge transfer.

4.5.3 Theoretical Framework and Research Issues

Cognitive Flexibility Theory

Cognitive Flexibility Theory (CFT) (Spiro et al., 1988) is a theory of learning and instruction that seeks to provide a framework for both understanding and promoting the advanced acquisition of knowledge, particularly in complex domains that possess a high degree of conceptual irregularity or ill-structuredness. (Due to space limitations, this section will provide only a brief discussion of CFT. For a fuller explication of this theory, see Spiro, Vispoel, Schmitz, Samarapungavan, and Boerger, 1987; Spiro et al., 1988; and Spiro and Jehng, 1990.) An important aspect of the theory, that is a central research issue in this study, is its concern with *knowledge transfer*.

An analysis of learning failures in several recent studies has identified various forms of *oversimplification of complex knowledge* as a significant contributing factor (Coulson, Feltovich, and Spiro, 1990; Feltovich et al., 1989; Myers, Feltovich, Coulson, Adami, and Spiro, 1990; Spiro, Feltovich, Coulson, and Anderson, 1989). These sources of failure in understanding are also believed to contribute to deficiencies in transfer ability. To address these problems, CFT makes a number of instructional prescriptions, such as: use of *multiple knowledge presentations* (e.g., multiple themes, analogies, case examples, lines of argument); *explicitly linking and tailoring concepts to practice or case examples* (i.e., situating conceptual knowledge in contexts that are similar to those required for the application of the knowledge); *incrementally introducing complexity* in small, cognitively manageable units; stressing the *interrelated and web-like nature of knowledge* (instead of isolated and compartmentalized knowledge); and encouraging the *assembly* of appropriate knowledge from various conceptual and case sources (rather than the intact retrieval of previously memorized information) (Spiro et al., 1988; Spiro and Jehng, 1990). The use of instructional features such as these are predicted by CFT to help the learner cultivate a multifaceted and flexible cognitive representation of complex knowledge. Such cognitive representations should enhance knowledge transfer by increasing the likelihood that an appropriate compilation of knowledge can be constructed to fit the understanding or problem-solving needs for the wide variety of case characteristics associated with ill-structured domains.

Research Issues

There is preliminary experimental evidence from two noncomputer studies that procedures based on Cognitive Flexibility Theory in fact promote the ability of students to acquire and transfer complex knowledge (Hartman and Spiro, 1989; Spiro et al., 1987). In

addition, CFT design prescriptions have also been used to develop several prototype instructional hypertext systems (Spiro et al., 1988; Spiro and Jehng, 1990). The present study was intended to extend this research and development program to specifically test selected CFT theoretical and design issues associated with the use of hypertext learning environments.

The main experimental hypothesis of this study was that a hypertext approach incorporating several design features based on CFT would promote the acquisition of complex material (the social impact of technology) in a manner that would enhance the learners' ability to apply the knowledge in cognitively demanding and novel knowledge application tasks. A control condition--a hybrid hypertext/computer-based drill--was developed which employed a number of antithetical design features that presented the same instructed knowledge in a more rigid and decontextualized manner. While the latter treatment was predicted to foster the efficient and effective acquisition of factual knowledge in the domain, it was expected that subjects in this group would perform at a lower level on tasks requiring knowledge transfer (i.e., the application of instructed knowledge to conditions unlike those found in the initial learning) than would the subjects receiving the experimental hypertext treatment.

4.5.4 Overview of Hypertext Study

Experimental Features

The social impact of technology was selected as the topic for the instructional content. This is a complex domain that is conceptually ill-structured due to the multidisciplinary focus of the literature and the rich and dynamic conditions that characterize real world applications of this knowledge. Thirty four subjects were randomly assigned to the two main treatment groups for a study that was administered in two hour sessions over a period of four days. An initial period of knowledge acquisition, referred to as *Reading Time*, provided an introductory exposure to the subject matter using similar hypertext materials read by both groups. A period of more advanced study was provided during *Study Time* in which the contrasting treatments for the two groups were administered. The main component of the treatment for the experimental group, referred to as *Thematic Criss-Crossing*, involved a hypertext environment in which abstract and case-centered components of the domain were studied from a variety of intellectual perspectives that highlighted different aspects of the knowledge. The control condition consisted of a computer-based drill that provided practice on factual items drawn from *Reading Time* and on items requiring the recognition of abstract knowledge components in previously read passages. The primary evaluation tasks, which were administered on the second and the fourth sessions, consisted of short answer tests on the acquisition content and written essays requiring the subjects to analyze and resolve issues posed in new social impact of technology situations.

Major Results

The short answer items were intended to measure the acquisition of factual knowledge associated with the instructional content. As predicted, the hybrid hypertext/computer-based drill group achieved significantly higher scores on this test of lower-level domain knowledge (means of 52.6 vs. 47.4, averaged across the two conditions). We feel that there were several aspects of the computer-based drill treatment that contributed to this finding, such as the *simplification* and *abstraction* of the domain knowledge resulting from studying the material

using specific multiple-choice questions on factual and theme content and the implicit emphasis on rote memory of the material resulting from recycling missed items for further practice in the drill format (see also Jacobson and Spiro, 1991a).

While there was no difference between the groups on the judges' scores for the more cognitively demanding Problem-Solving Essay task administered at the end of Session 2, the Thematic Criss-Crossing group achieved significantly higher essay scores by Session 4 (9.5 vs. 7.2 after Session 4; after Session 2 the nonsignificant difference was 8.4 vs. 7.6, in the same direction). The delayed effectiveness of the Thematic Criss-Crossing treatment was expected as it was thought that a critical period of time would be required by the subjects to assimilate the multiple and interrelated perspectives of the knowledge presented by this treatment. Subjects receiving the Thematic Criss-Crossing treatment thus were better able to apply their acquired knowledge of the social impact of technology in an unfamiliar within-domain problem-oriented situation.

4.5.5 Hypertext and Learning Complex Knowledge for Transfer: Theoretical and Design Considerations

A short duration study such as this (over a two week period with about four hours of instruction) is but a snap-shot of advanced learning in a complex and ill-structured knowledge domain. It is therefore notable that significant experimental effects were achieved in the limited amount of time available for the study that were highly consistent with the study's cognitive theoretical framework. This section will examine the experimental findings from the perspective of five Cognitive Flexibility Theory instructional prescriptions (discussed earlier) that were incorporated into the hypertext system used in this study. The experimental Thematic Criss-Crossing group received all five of these prescriptions while the hybrid hypertext/computer-based drill received only three of the prescriptions, along with a computer-based drill.

Cognitive Flexibility Hypertext Features Common to Both Experimental and Control Groups

Use multiple representations of knowledge. A central recommendation of CFT concerns the use of multiple knowledge representations to help avoid the oversimplification of complex knowledge. The hypertext materials used during Reading Time were intended to represent the complexity of the social impact of technology domain through the use of multiple case examples and multiple thematic or conceptual perspectives. Six cases were written as the primary hypertext knowledge nodes that described different examples of the social impact of technology, such as a case describing an attempt to develop an electric car in France during the 1970s and a case discussing the use of human subjects for medical research on the development of an artificial heart. Multiple conceptual perspectives were provided through the use of six broad-scope and partially overlapping thematic dimensions of analysis that were identified from the literature, such as *Freedom - Control* and *Community - Alienation*.

Link concepts to case examples. The six cases were subdivided into *minicases* that served as the primary instructional units/hypertext nodes. Each minicase consisted of a short section of text describing a portion of the case, a list of relevant themes, and theme commentaries (see Spiro and Jehng, 1990). The text of each minicase was analyzed to determine which of the six

social impact of technology themes were relevant; applicable themes were placed in a Theme List node for the subject to read. Theme commentaries were written that were intended to be representative of expert knowledge in this domain. The commentaries were specific to every minicase for each of the listed themes. The theme commentaries explicitly linked the generic definitions of the thematic perspectives (i.e., abstract knowledge components) to the specific features of text sections used in the minicases (i.e., case-centered knowledge components).

Incrementally introduce complexity. The minicase sections represented small and cognitively manageable units of study that still reflected aspects of the domain's conceptual complexity and ill-structuredness even at the beginning of the instructional treatment. The instructional sequence provided increasing amounts of complexity as the subject read more social impact of technology cases. Also, because of the different combination of themes associated with each minicase, variability of conceptual or thematic coverage could be seen both within and between cases. Concepts were thus *not* presented in an oversimplified or abstracted manner initially as in many curricular approaches, with the potential problem that the early simplifications would interfere with learning domain complexities at a later phase of instruction (Feltovich et al., 1989; Spiro et al., 1987; Spiro et al., 1988, 1989). Instead, the early instructional sequence reflected a moderate degree of the richness associated with the domain that was then elaborated on with later Study Time learning activities.

Cognitive Flexibility Hypertext Features of the Experimental Thematic Criss-Crossing Group

This section discusses two additional tenets of Cognitive Flexibility Theory, the *web-like nature of knowledge* and *knowledge assembly from different conceptual and case sources*, that were implemented in the Thematic Criss-Crossing treatment administered to the experimental group during Study Time. It is important to note that while these two additional prescriptions were added during Study Time, Thematic Criss-Crossing continued to utilize (in slightly different ways) the three prescriptions discussed earlier. The *multiple representations of knowledge* and *incrementally introducing complexity* prescriptions were instantiated through the use of multiple combinations of themes from which to explore the minicase knowledge base in novel and intricate ways (i.e., thematic criss-crossing), and the *linking of concepts to case examples* was continued in a similar manner as in Reading Time through the use of the same theme lists and thematic commentaries (although subjects were required to read only selected theme commentaries associated with a particular Thematic Criss-Crossing combination). The next two sections briefly discuss the additional CFT prescriptions employed in the Thematic Criss-Crossing treatment.

Stress interrelated and web-like nature of knowledge. As noted earlier, Thematic Criss-Crossing required the subjects to reread minicases that exemplified different combinations of themes. It was hoped that this procedure would help demonstrate case-based conceptual interrelationships in the domain knowledge that may have been unnoticed when initially studying the material in Reading Time. For example, reading minicase A with the themes "x y z," minicase B with themes "u x y," and minicase C with themes "v z" may not suggest any underlying knowledge interrelationships to a novice in the domain. However, considering themes "x y" together demonstrates knowledge linkages between minicases A and B, while a consideration of theme "z" illustrates an overlap between minicases A and C. Each combination of themes allows a different traversal of the minicases that comprise the knowledge base of the

hypertext. Such a theme-based exploration of the minicases is intended to highlight many different facets of the domain knowledge. In all, the Thematic Criss-Crossing subjects read 15 different thematic combinations over the four experimental sessions to implement the CFT prescription to demonstrate the web-like structure of the conceptual and case interrelations in the instructional domain.

Encourage assembly of knowledge from different conceptual and case sources. The Thematic Criss-Crossing treatment was designed to provide an externalized, technology-based procedure that explicitly demonstrated how knowledge could be assembled from different conceptual and case sources, a procedure that could be considered analogous to knowledge representation schemes observed in experts. For example, protocols of experts engaged in problem-solving tasks in a variety of domains have documented their ability to flexibly apply knowledge and solve problems in a domain (e.g., Chi, Glaser, and Farr, 1988). It may be that the *assembly* of knowledge from various conceptual and case sources rather than just the intact retrieval from memory of precompiled knowledge structures (Spiro et al., 1988) is a component underlying the flexible knowledge application and problem solving performance observed in experts. Part of the effectiveness of the Thematic Criss-Crossing treatment thus may be due to the externalized *modeling of knowledge assembly* demonstrated through the computer re-editions of the link relationships between the minicases based on different theme combinations. With this procedure, the student directly experiences the combination and recombination of knowledge components that, internalized over time, undergirds the development of cognitive flexibility and the improved ability to transfer knowledge to new situations.

Implications for the Design of Hypertext Learning Environments

There are important implications for the design and use of instructional hypertext systems that are suggested by the findings of significant transfer effects associated with the experimental Thematic Criss-Crossing treatment on the Problem-Solving Essay task. While a full consideration of these issues is beyond the scope of this chapter, a critical implication of the study is that *merely providing hypertext links to different information nodes is not sufficient to promote the transfer of knowledge to new situations*. Both groups had an equal opportunity to learn abstract and case-centered components of the domain knowledge with the hypertext materials used during Reading Time and both groups also received further study in the domain during the contrasting Study Time treatments. As noted earlier, the superior performance of the Thematic Criss-Crossing group on the Problem-Solving Essay suggests that the ability to acquire and transfer complex domain knowledge requires *special design parameters* for the hypertext learning environment, such as presenting the knowledge in multiple and interconnected ways. Future research is obviously needed to study the effects of such theory-based design parameters on the effectiveness of hypertext learning environments in a variety of content areas and instructional situations.

4.5.6 Conclusion

This chapter discussed a number of theory-based features of an approach to structuring a hypertext system that was used in a recent study of learning and knowledge transfer in a complex and ill-structured content area. Significant transfer effects were found for the experimental hypertext treatment that employed recent cognitive instructional theory

prescriptions such as the use of *multiple knowledge representations*, depicting the *interconnected and web-like* structure of abstract and case-centered knowledge components, and stressing the *assembly* of knowledge from different conceptual and case sources. The rigorous evaluation of instructional hypertext design principles is still a relatively new area of cognitive research. Given the proper attention to critical theoretical and empirical issues associated with learning complex knowledge, hypertext learning environments show promise of being a valuable instructional supplement to other types of cognitively oriented computer-based learning systems such as microworlds, intelligent tutoring systems, and empowering environments.

4.5.7 Footnote for Chapter

The principal contributors to this chapter were Michael Jacobson and Rand Spiro.

**Chapter 4.6 A Detailed Example of the Application of Cognitive Flexibility
Theory to a Hypertext System for Case-Based Development of
Conceptual Understanding in an Ill-Structured Domain:
The Military Strategy of the Indirect Approach**

4.6.1 Introduction

The preceding chapter presented the theory of Cognitive Flexibility Hypertexts in extensive breadth and detail. The present chapter is intended as a much more concrete follow-up to that theoretical presentation, more closely approximating the way the system would actually be used in the initial stages of learning.

Purpose of the Chapter

There were several purposes for the research discussed in this section. First, the research was carried out in an archetypal arena for relating *conceptual knowledge* to *decision-making*: an ill-structured domain of *strategic action*. In particular, we endeavored to set up a cognitive flexibility hypertext system to teach Liddell-Hart's military *strategy of the indirect approach* (1954), as applied to actual military battles, based on a Cognitive Flexibility Theory model of advanced knowledge acquisition for the domain. Second, the domain chosen, the military strategy of the indirect approach was chosen for its direct and obvious implications for the military.

Third, the research presented in this section is intended to provide a fairly detailed working out of the theoretical notions discussed in the earlier sections. That is, it was intended to be an *instantiation* of Cognitive Flexibility Theory, both at the cognitive and instructional levels. Many of the elements of Cognitive Flexibility Theory discussed in different places earlier in this Report will be illustrated in this concrete case; thus, this section serves a concrete *integrative* function. Tenets of Cognitive Flexibility Theory will be highlighted at their points of instantiation to facilitate this process of integration.

Fourth, an important point we want to make is that learning approaches of the *situated cognition* type (e.g., Brown, Collins & Duguid, 1989) have had two important shortcomings (along with their many strengths and positive contributions): 1) they have focused on the importance of the situation or case *at the expense of conceptual knowledge* -- Cognitive Flexibility Theory emphasizes the importance of *joint* relationships among the two; and 2) they have not presented specific models of situated learning systems and processes. The work discussed in this system is intended as a *worked out model of one variety of situated learning, incorporating a central role for conceptual information as knowledge-in-action*. The point of Cognitive Flexibility Theory is that neither conceptual understanding nor case-based understanding can be dominant. Rather, a new view of conceptual understanding and its acquisition is required for ill-structured domains, a view that has at its core the notion that the meaning of conceptual entities is at best incompletely expressed in the abstract, and full meaning is possible only in a particular situation of use. This new view is explicated and exemplified in this section of the Report.

The plan for this chapter is as follows. First, the the Strategy of the Indirect Approach is described. Then, an overview is presented of how a general plan for learning the Strategy of the Indirect Approach is derived from Cognitive Flexibility Theory. The bulk of the rest of the paper is devoted to a *tour* of a Cognitive Flexibility Hypertext system for teaching the Strategy of the Indirect Approach, with examples drawn from the *early stages of learning* using the system, using actual historical cases as a basis for instruction. Throughout the discussion ways that the Cognitive Flexibility Hypertext follows the instructional tenets of Cognitive Flexibility Theory will be highlighted.

Introduction to the "Strategy of the Indirect Approach"

The present chapter illustrates the application of Cognitive Flexibility Theory to a Cognitive Flexibility Hypertext for the military strategy of the indirect approach (SIA), derived from the thinking of Liddel-Hart (1954), as updated by contemporary military strategists of the *operational level* of war (e.g., Luttwak, 1986). The strategy of the indirect approach is crucial to the application of such contemporarily central strategies as relational maneuver and blitzkrieg. SIA requires commanders (in headquarters and in the battlefield) to flexibly adapt to dynamically changing battle situations and to most efficiently leverage their own forces by striking at weak points of the enemy -- all to produce the maximum psychological dislocation of the enemy. (A more detailed understanding of what SIA is will emerge during the course of this chapter, as we explore a Cognitive Flexibility Hypertext that is intended to teach it.)

SIA is likely to continue to be the dominant strategy of the current era, given limitations of resources and uncertainty as to the characteristics of the enemy and the arena of conflict as the military must be ready to deploy virtually anywhere in the world. SIA was the strategy we employed in the recent Persian Gulf War. Unfortunately, SIA has often been thought by military experts to be nonteachable because it relies too much on the "art" of command or the "genius" of individual commanders. We argue that the difficulty in teaching this strategy is due to its being treated conventionally, i.e., as a well-structured domain, when in fact it is highly irregular and ill-structured, and hence requires an approach that stresses cognitive flexibility.

Aspects of Knowledge Structure and Its Acquisition in an Ill-Structured Domain

In this section of this chapter, we will illustrate some of the crucial features of the structure of knowledge for this action-oriented strategic approach. Since the Cognitive Flexibility Hypertext for the Strategy of the Indirect Approach is derived directly from the learning implications of these special knowledge structural features, the nature of the hypertext approach will be naturally introduced in the course of the analysis.

What is the Strategy of the Indirect Approach? Does it have a finite, denotatively specifiable *meaning*? The answer is "no", as it would be for any ill-structured conceptual topic (e.g., Spiro et al., 1988; Wittgenstein, 1953). How, then, can the topic be represented and, relatedly, taught? The first step is to acknowledge its complexity: SIA has been thought of as 'essentially' many different things, a theory of flexibility, of psychological and physical dislocation, of indirection, of momentum, of expediency, of distraction and deception. And all of these attributions are true, although SIA is not 'essentially' any of them. In different cases, one

or the other of these features will appear to dominate the others in its importance for grasping SIA. However, shifts in dominance patterns for other cases are always easy to demonstrate. Furthermore, none of these features of SIA is present in every case of the approach's application. Perhaps most importantly, none of the features of SIA can itself be denotatively defined, but rather is characterized by a relationship among its instances of rough family resemblance (without any exemplary instances that could adequately serve as prototypes for the entire set). Thus, the conceptual entities that make up SIA are themselves ill-structured, just as is the domain as a whole.

Therefore, following the dictates of Cognitive Flexibility Theory, the conceptual entities of SIA (e.g., *flexibility*, *dislocation*) must be taught through their *application or use* in actual battle situations, treated as no more than *themes* that run through cases in a less than orderly and regular manner, and dealt with in combination rather than isolation. The family resemblance patterns that unite instances of each SIA theme, as well as the patterns that produce usable resemblances across instances and non-instances of the overall SIA approach, are gradually accumulated from structured experiences with their occurrences across actual cases. This learning process, and how it relates to the use of SIA knowledge in making decisions in new cases, is the subject of the remainder of this chapter.

Disclaimer Concerning Limited Objectives

It is important to understand the limitations of the SIA Cognitive Flexibility Hypertext. At this stage, we are *not* trying to prepare people for decision-making. Rather, the goal is to teach the conceptual structure of SIA, an ill-structured domain of concepts-in-action, as it is applied to actual historical battles (i.e., as a system integrally connected to application in real cases, not as an abstract entity in isolation from application). This is a large step towards full-fledged decision-making, however.

A related point is that we are not arguing that studying our vicarious cases is a *substitute* for actual participation in "live" cases. Rather, we argue that the two case approaches are *complementary*. Real-time participation in cases has the limitation that they take too much time, and therefore you can not cover enough of them in an instructionally manageable amount of time to acquire sufficient experience with the variety of structural forms the cases take in an ill-structured domain -- this is one of the reasons that professionals usually require so many years of experience (see Chapter 4.4). You need *many* cases for sufficient experience for transfer in ill-structured domains. However, we always use *real* cases (rather than artificially constructed ones) in our Cognitive Flexibility Hypertexts -- you have to get people ready for what they will have to apply their knowledge to. All of this is not to say however, that we eschew active participation; in fact, the latter is an important tenet of Cognitive Flexibility Theory and is part of the Cognitive Flexibility Hypertext approach to teaching SIA. It will not be showcased in the present chapter, however. Here we just stress the aspect of the Cognitive Flexibility Hypertext concerned with reading to learn conceptual structural relations as they relate to cases--a new kind of case-based *conceptual* instruction.

Also, in the limited space available, we will provide only a limited number of "criss-crossings." We just want to show how much *manageable complexity* you can get from just a

couple of hours or so of instruction. (Cognitive Flexibility Theory says you then get an exponential increase with additional time spent.)

Also keep in mind that a lot of what you will read about SIA will not be easy. But that is the way that it has to be: SIA is hard, and although we believe we make it as easy as possible using the Cognitive Flexibility Hypertext approach, it would not do to make it so easy that you could not learn the difficult topic that it is your goal to learn. Furthermore, in actually using the program, you would not try to learn everything at one time, because you would know that you would be revisiting the same information later. This would remove some of the pressure of achieving a full understanding the first time or two that you read material.

In sum, before the reader embarks on this little exploration of the SIA Cognitive Flexibility Hypertext it is recommended that he or she keep two things in mind: 1) it is only a demonstration of how Cognitive Flexibility Hypertexts do *one* thing, *conceptual knowledge structure building* -- there are lots of important things in the theory of Cognitive Flexibility Hypertexts that are not addressed in what follows (e.g., active participation; real-time situated learning); and 2) it is only a *little* demonstration, barely scratching the surface of what could be done in many more "criss-crosses," with many more cases. It is just an introductory *tour* through a Cognitive Flexibility Hypertext in the initial stages of learning SIA.

4.5.2 Cases in Cognitive Flexibility Hypertexts: The 1940 Ardennes Campaign as a Staging Ground for SIA Instruction

A central tenet of Cognitive Flexibility Theory is that learning in ill-structured domains must be *case-centered* (because of the differences across cases of any possible classificatory type). For the purposes of this demonstration of the nature of Cognitive Flexibility Hypertexts and how they relate to the foregoing work on failures in advanced knowledge acquisition and associated remedies in Cognitive Flexibility Theory, we will present one case in detail, the German westward advance through Belgium in 1940, emphasizing the surprise breakthrough in the Ardennes Mountains and its aftermath, culminating in the evacuation of Allied forces at Dunkirk.

We begin with a first pass introduction to SIA and a description of the campaign, like that which would be presented to a learner beginning instruction in SIA. Note that the campaign has been broken up into *nine* chronologically based segments or *scenes*. As we will see, these scenes are used as the unit of instruction, not the campaign as a whole. This is a crucial *design decision* in Cognitive Flexibility Hypertexts, related to the problems of moving from simple to complex learning discussed in various places earlier and to receive more attention later.

Instructions to the reader: You should now read the introduction just below, and then proceed to read the Ardennes scene descriptions numbered 26-34 in conjunction with looking at the copy of the computer display of the battle map for each of the scenes beginning on the second page following (the scene number is in the upper left-hand corner). Again, be sure to consult the map for each scene in conjunction with the corresponding verbal scene description. The verbal descriptions are what you would see if you clicked on the "Description" button at the bottom of the figures with the maps.

Introduction to SIA and the Ardennes Campaign

The purpose of this learning experience is to teach you about a particular kind of strategy called the STRATEGY OF THE INDIRECT APPROACH (which we will usually refer to as SIA). The strategy can be applied in many different kinds of situations. However, it has been used most often in the context of decision-making for military battles, so we will teach the strategy in that context. After learning about the strategy, you will be given some test questions that ask you to apply the strategy to new situations, some military and some non-military.

The STRATEGY OF THE INDIRECT APPROACH (SIA) is, in a sense, paradoxical. It is a strategy that says you should not always take the most direct and obvious approach to attaining some objective, that you should not necessarily take the path of least effort, travel the shortest distance, or follow expected "rules" of action. But you should be on your guard -- SIA is a complex strategy, with many facets; learning SIA will require much more than just memorizing a simple definition (like the one just presented).

We will begin by presenting a historical example that will be used to illustrate SIA. After the example is fully presented, the concepts that make up SIA will be introduced and applied in the context of that example. The example is the Ardennes offensive in 1940, in the early stages of World War II; up to the very end of the campaign, the Germans exemplify SIA while the Allies violate SIA consistently; in the last part of the campaign the roles are reversed.

REMEMBER: THE MAPS ON EACH OF THE NEXT NINE PAGES CORRESPOND TO SCENES IN THE ARDENNES CAMPAIGN. DESCRIPTIONS OF WHAT IS GOING ON IN EACH OF THE SCENES IS TO BE FOUND ON THE PAGES IMMEDIATELY FOLLOWING THE LAST MAP. READ THE DESCRIPTIONS IN CONJUNCTION WITH LOOKING AT THE MAP. [Of course, in the hypertext you would not have to page back and forth to accomplish this.]

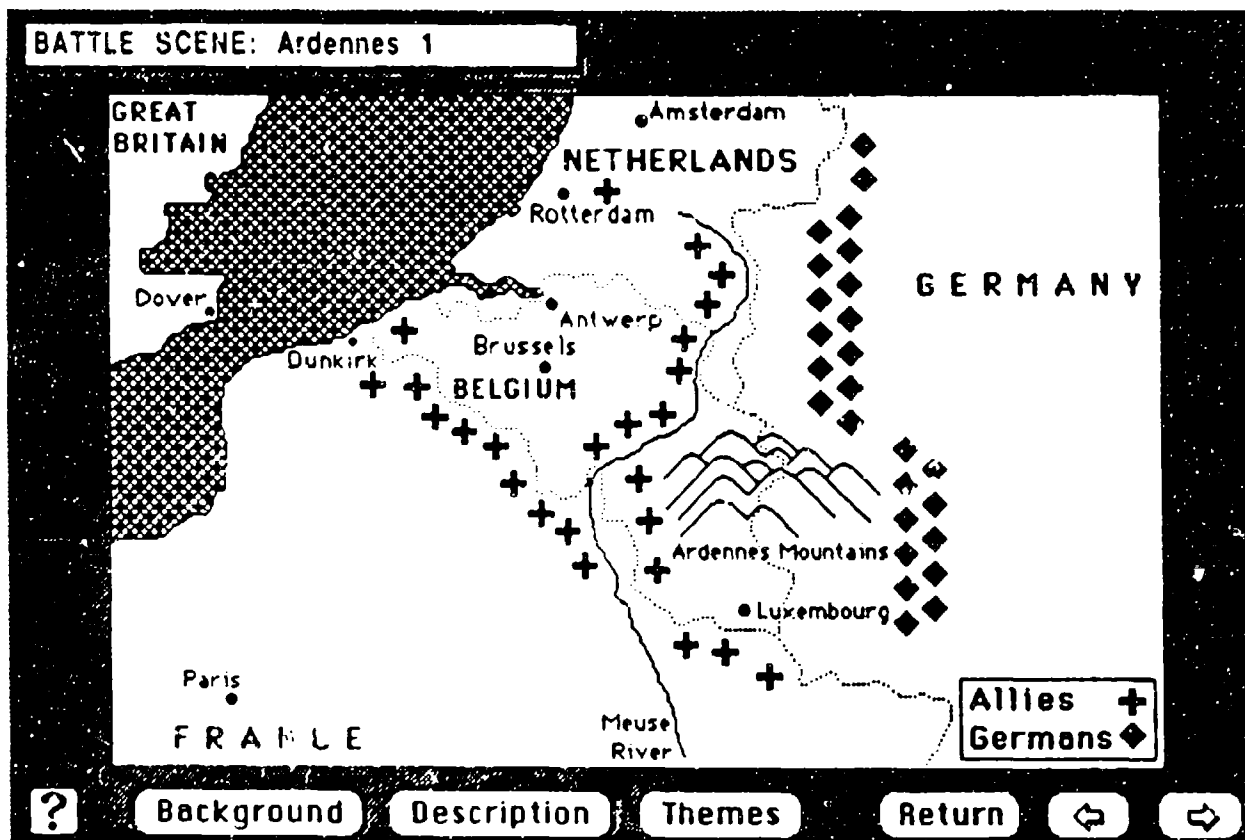


Figure 2f. Scene 1

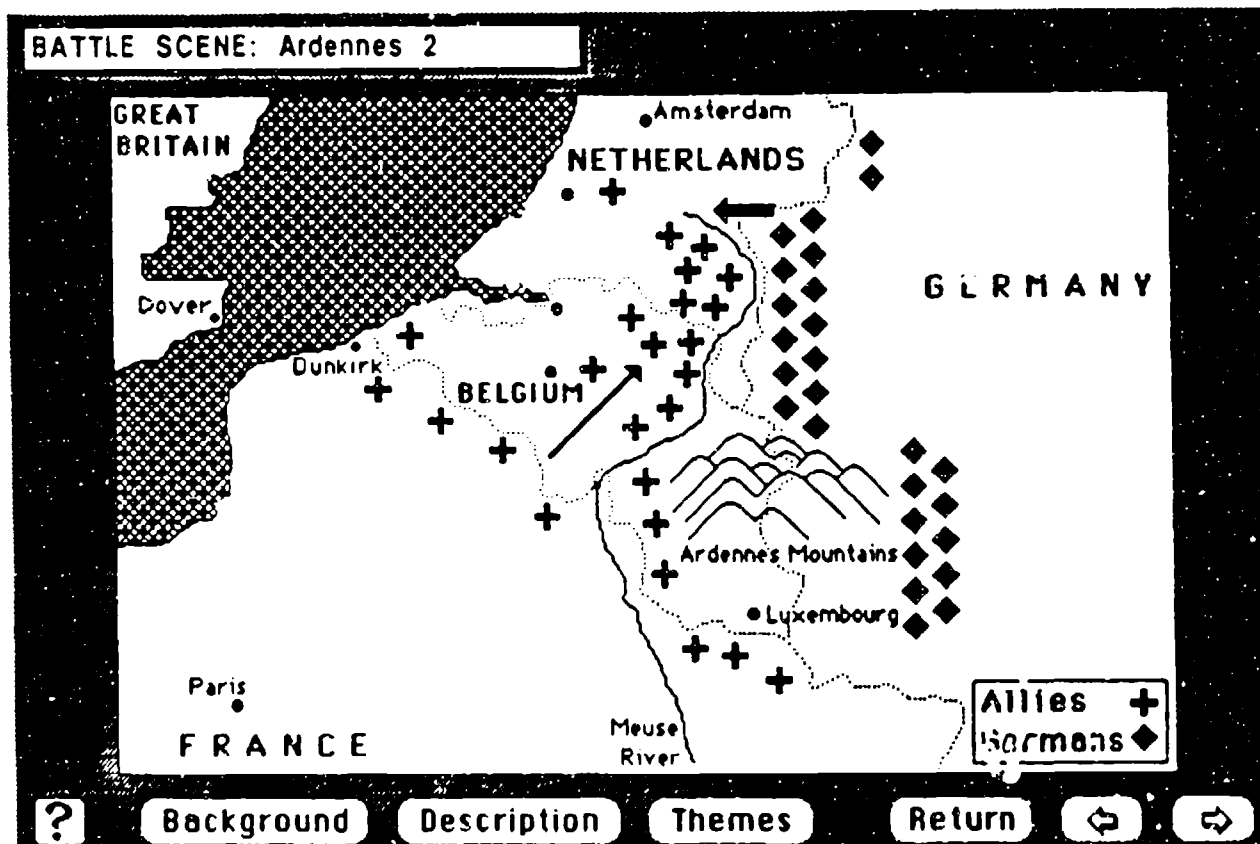


Figure 27. Scene 2

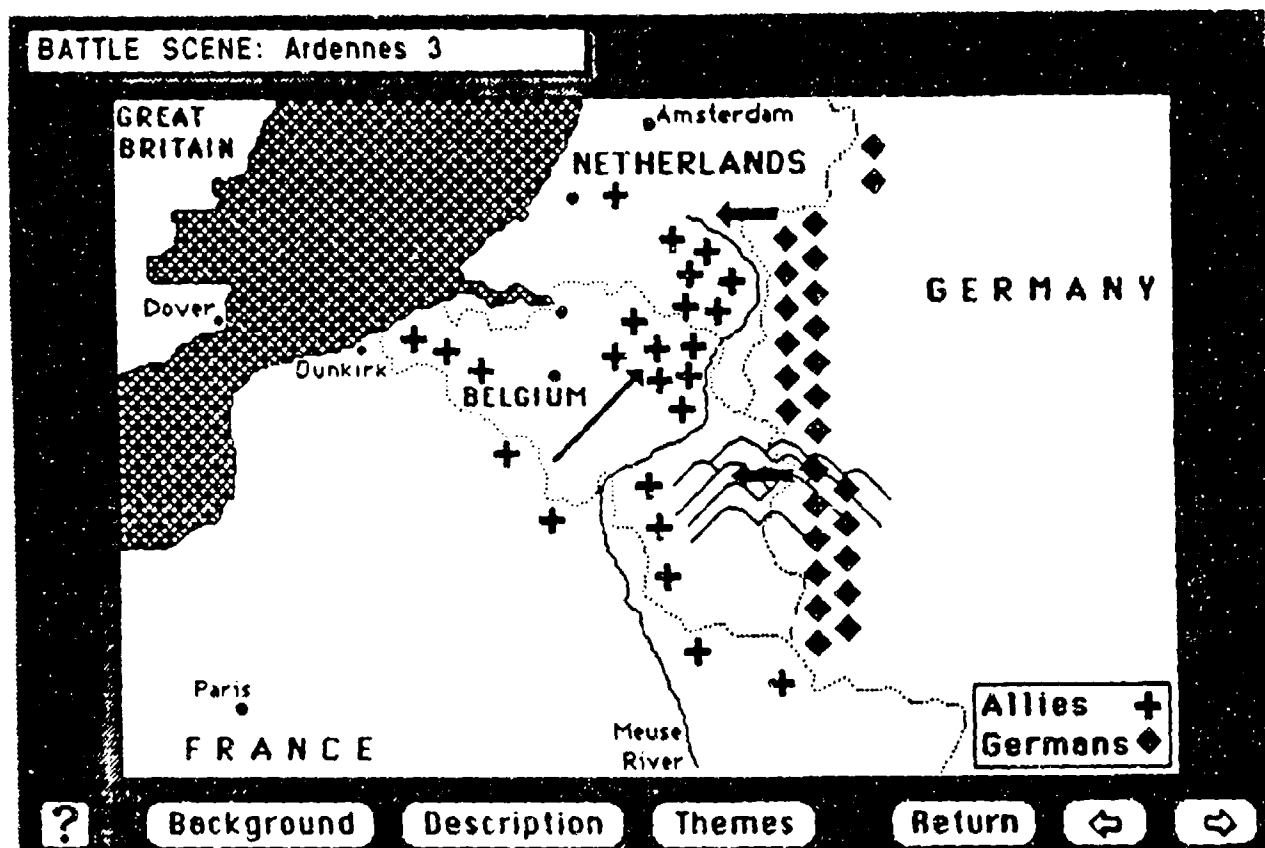


Figure 28. Scene 3

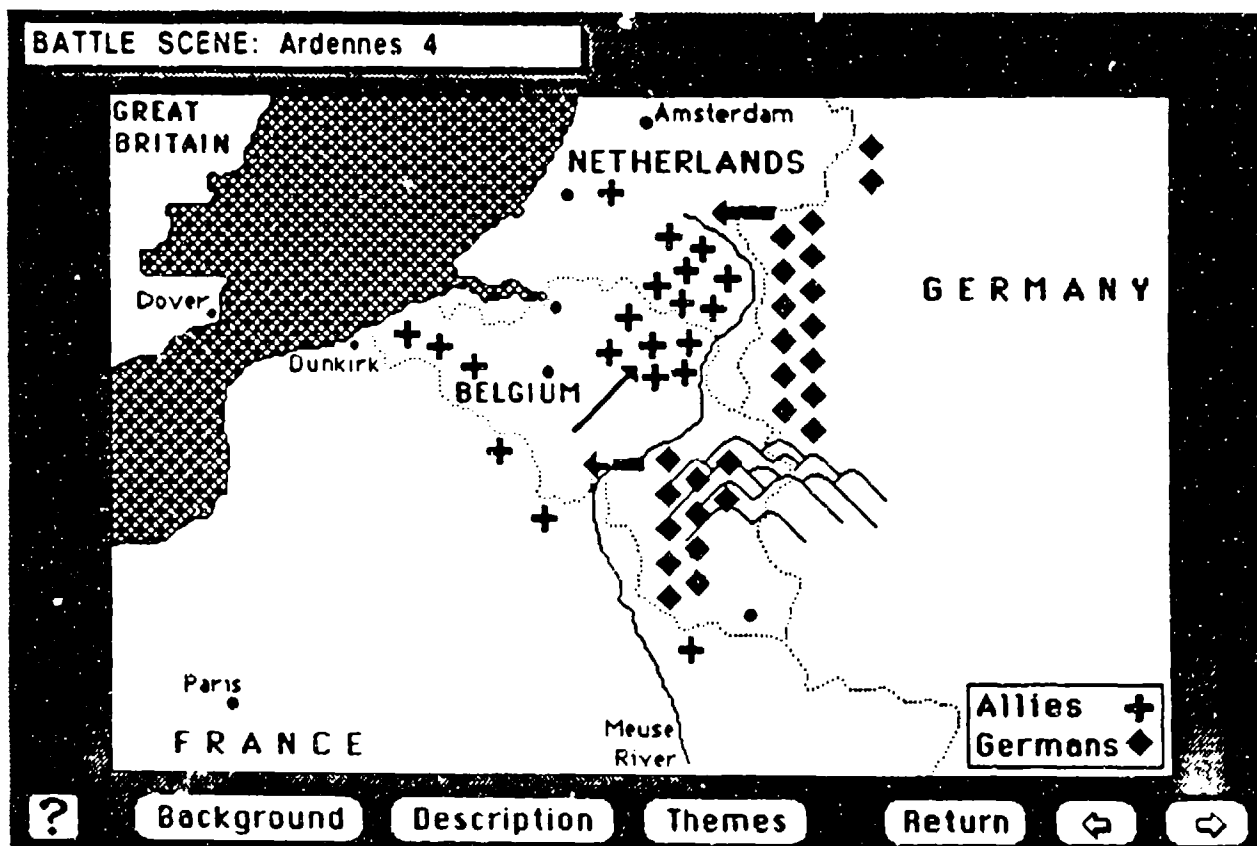


Figure 29. Scene 4

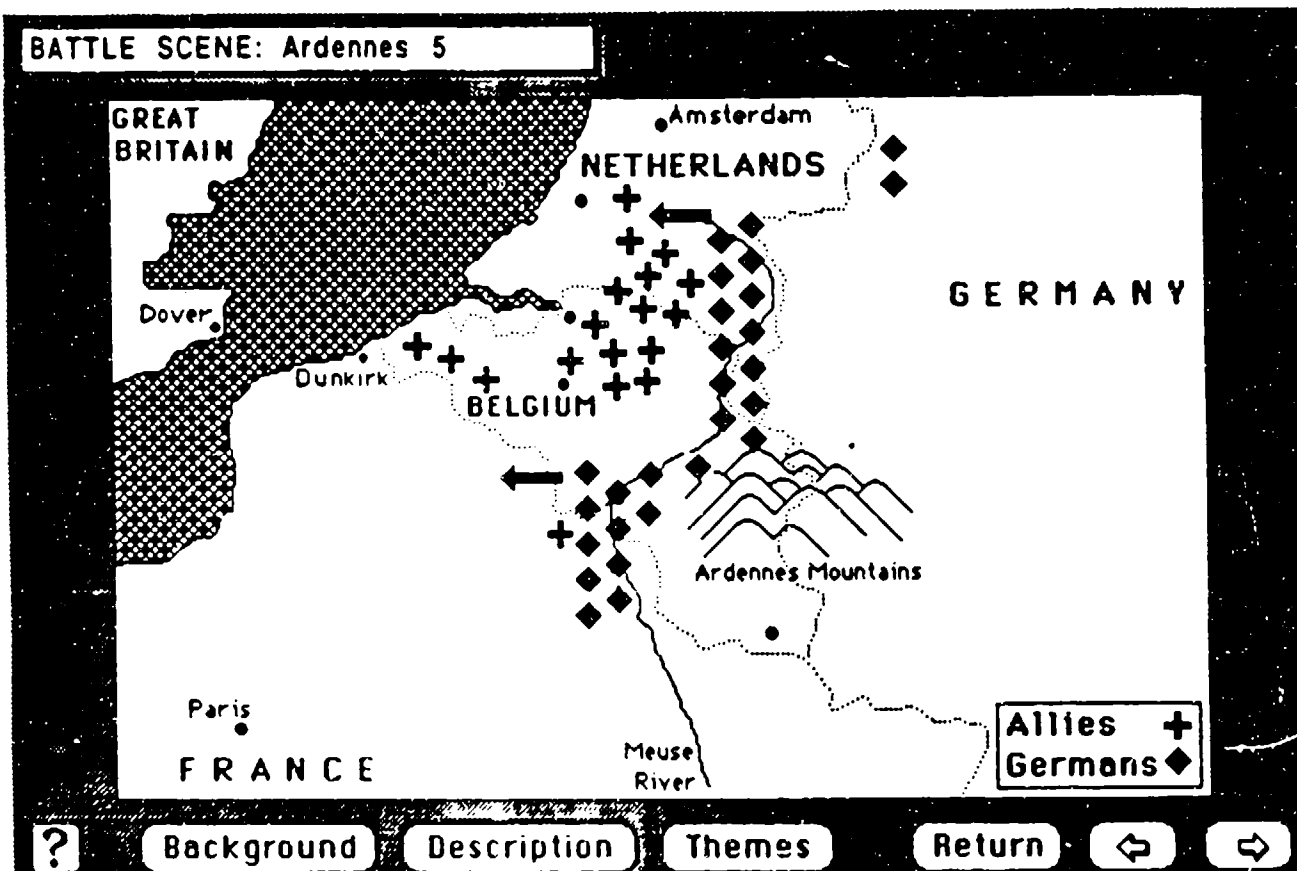


Figure 30. Scene 5

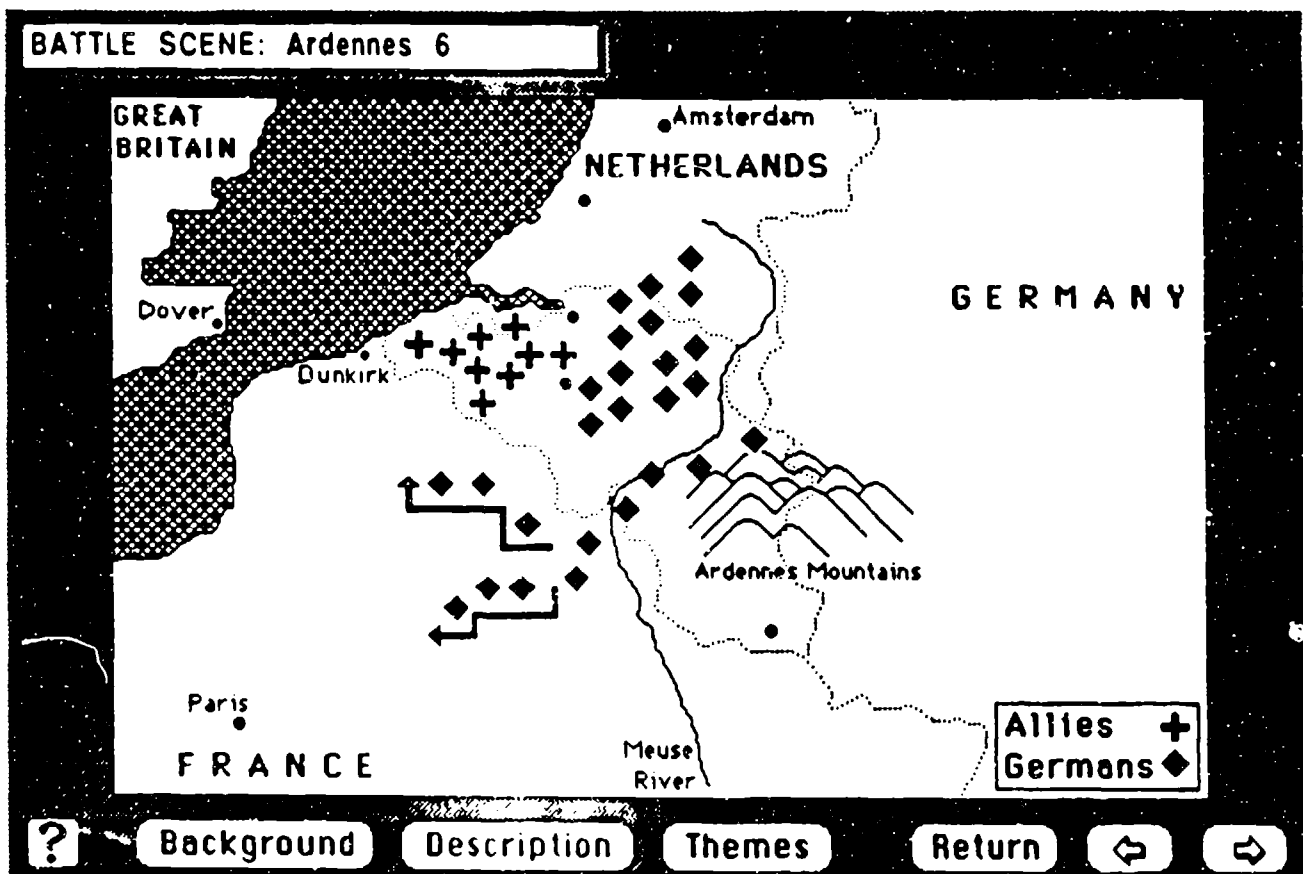


Figure 31. Scene 6

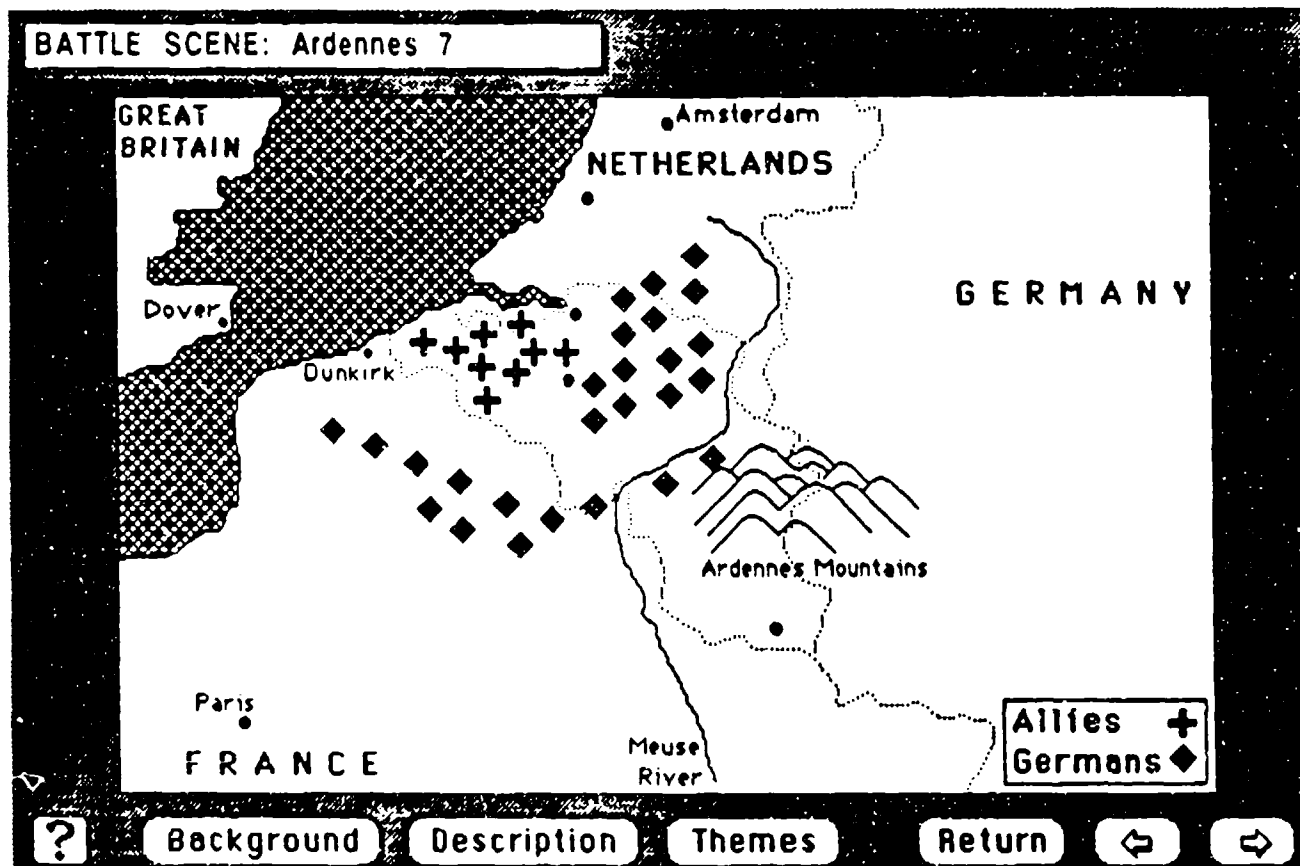


Figure 32. Scene 7

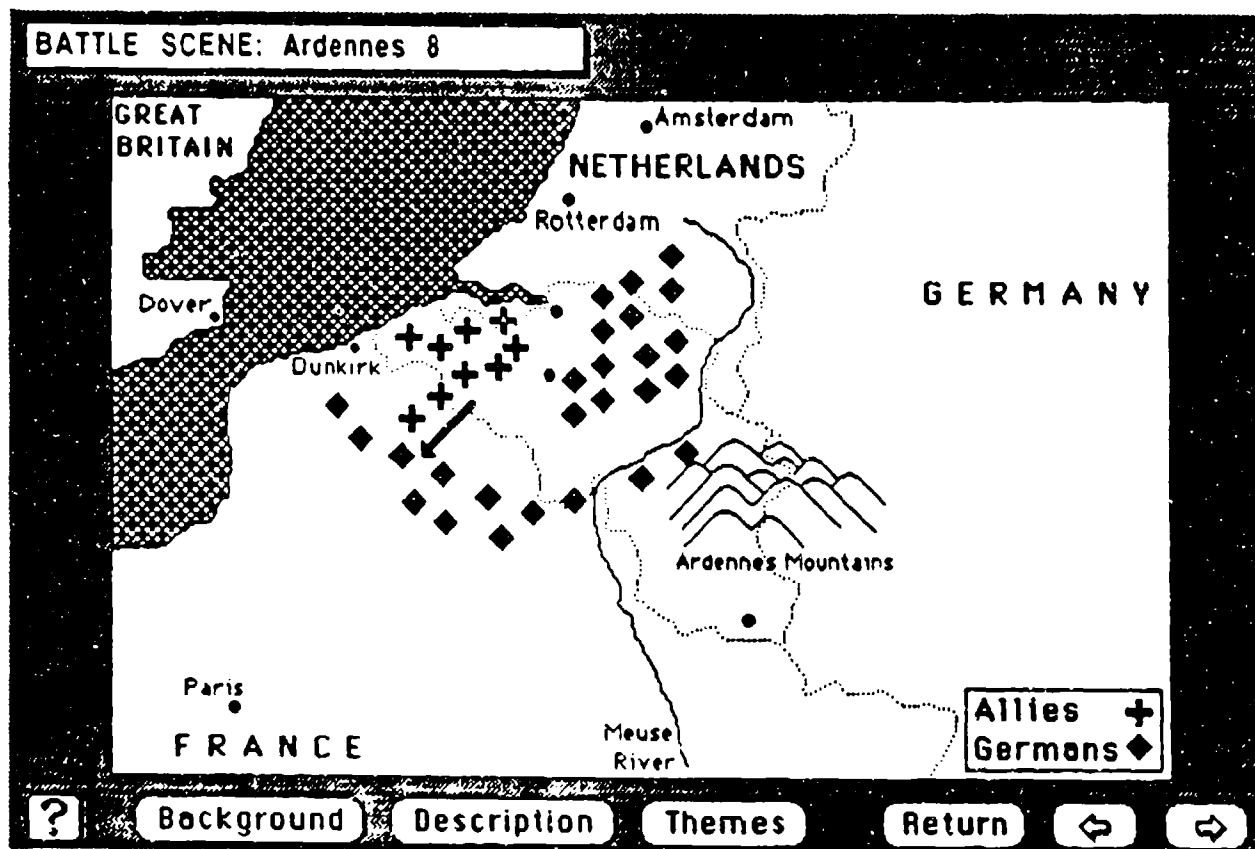


Figure 33. Scene 8

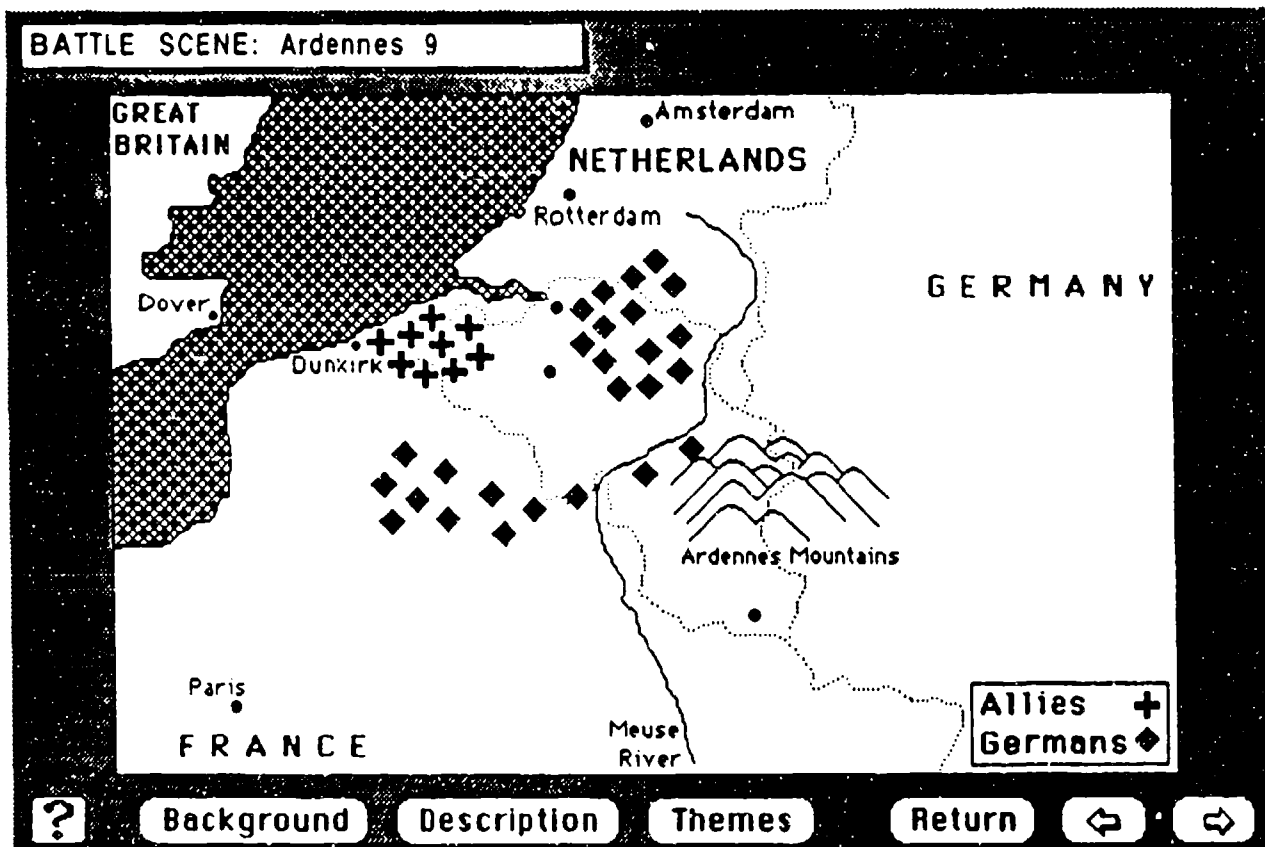


Figure 34. Scene 9

Scene Descriptions for the Ardennes Case
(see Maps for each scene on the preceding 9 pages, Figures 26-34)

BATTLE SCENE: Ardennes 1

BATTLE BACKGROUND

Allied forces anticipate German attack through the central Belgian plain, and have planned accordingly.

BATTLE DESCRIPTION

German feints SHAPE' OPPONENT to elicit Allied responses that signal the Allies' intention to advance rapidly into central Belgium.

BATTLE SCENE: Ardennes 2

BATTLE BACKGROUND

The Allies have a very strong expectation of a German offensive through the central Belgian plain.

BATTLE DESCRIPTION

A sizable German force attacks through central Belgium, as anticipated. Allied forces immediately stream towards the area of the attack in massive numbers, with the most mobile forces leading the way.

BATTLE SCENE: Ardennes 3

BATTLE BACKGROUND

An anticipated German invasion of central Belgium has led to a massive rapid movement of Allied Forces to that region.

BATTLE DESCRIPTION

A German offensive force is reported in the Ardennes Mountains region (thought to be impassable by large armoured forces). The Allies ignore these reports and continue to stream with full commitment into central Belgium, north of the Ardennes.

BATTLE SCENE: Ardennes 4

BATTLE BACKGROUND

The Allies have been responding en masse to an anticipated German attack in central Belgium. Reports of German forces in the Ardennes Mountains to the south have not been taken seriously.

BATTLE DESCRIPTION

It becomes clear that the Germans have launched a major offensive through the lightly guarded Ardennes region. This is particularly troubling because the overcommitted Allied forces are way out of position to react to the German move to the south (the most mobile Allied forces, having been sent into Belgium first, are the furthest in the wrong direction).

Also, the attack through the Ardennes simultaneously threatens: 1) the flank of the shaken Allied forces in Belgium, as well as their communications; 2) encirclement by a westerly followed by a northerly thrust; and 3) the vulnerable French capital of Paris.

BATTLE SCENE: Ardennes 5

BATTLE BACKGROUND

The Germans have broken through the lightly defended Ardennes in force, threatening multiple objectives (the flank of Allied forces in Belgium; encirclement of those forces; and vulnerable Paris).

The Allies are caught off balance from having thrust forcefully in the wrong direction and are out of position to counter the German move. They are shaken by the surprising and compromising turn of events.

BATTLE DESCRIPTION

The Germans move rapidly after breaking through at the Ardennes. They reach the Meuse River. Standard military procedure would be to wait for infantry and artillery to arrive and consolidate the position before crossing the river.

The German commander, placing a premium on the continuation of a rapid advance so as to further dislocate the Allies and to avoid momentum-halting contact, pushes immediately and successfully beyond the Meuse.

The Allies are still unable to respond. They are far out of position and off balance, uncertain where the Germans are going, and psychologically shaken.

BATTLE SCENE: Ardennes 6

BATTLE BACKGROUND

The Germans have been advancing rapidly and virtually unopposed following the Ardennes breakthrough. The Allies have been unable to respond.

BATTLE DESCRIPTION

The German army disperses somewhat and begins to follow a zig-zag, unpredictable course.

The Allies continue to be unable to make a credible response.

BATTLE SCENE: Ardennes 7

BATTLE BACKGROUND

The Germans have continued to perplex the Allies by pursuing an erratic direction of advance between a few alternative objectives.

The Allies have been unable to respond.

BATTLE DESCRIPTION

The Germans reduce their dispersion and zig-zagging, signaling their intention to move north and entrap the Allied forces between the southern army and the army advancing westward through the low countries.

BATTLE SCENE: Ardennes 8

BATTLE BACKGROUND

Up to the point of this scene, the Germans have been advancing boldly and quickly, confusing the Allies about their objective.

In the preceding scene the Germans had finally signaled their intention to swing northward.

BATTLE DESCRIPTION

The Allies launch a small counterattack at Arras. The Germans, not knowing the size of the attack and fearing their exposed flanks, halt their forward advance.

BATTLE SCENE: Ardennes 9

BATTLE BACKGROUND

The Allies have finally broken the German momentum, capitalizing on knowing where the German army would be concentrating to launch a small but very dislocating counterattack.

In response, the Germans have halted their advance to consolidate their position and defend their flank.

BATTLE DESCRIPTION

As the Germans pause to consolidate, the Allies are given important time to recover their balance and plan for the extrication of their trapped forces. The evacuation at Dunkirk followed.

Some Examples of Capabilities of Cognitive Flexibility Hypertexts

Cognitive Flexibility Hypertexts permit a number of kinds of manipulations of the instructional content. One set of options is operative for the *individual scene* (which we will sometimes call a *minicase*), and a second kind operates across scenes (within or across cases).

Scene-based (minicase-based) options. Look back at any of the Ardennes scene frames just presented. Several options are indicated along the bottom.

A mouse-click on "Description" provides a raw account of the events of the case with very little interpretation in terms of SIA. The descriptions of each scene in the program were presented before the immediately preceding figures. How it would appear when actually using the program is demonstrated in the first figure below (pressing the mouse at the downward arrow on the right of the text would cause the text to scroll down until it is completed).

The "Background" option, when selected, provides a description of the chronologically antecedent scene. (See the Figure 35 for an example.) This is important because of the nonlinear "criss-crossings" fundamental to Cognitive Flexibility Theory-based instruction, although the Ardennes case will have already been covered, it is still possible that moving to any single scene within the case, out of sequence, may temporarily disorient the learner. The "Background" information sets the stage for the current scene's exploration by *reminding* the learner about what was happening just before.

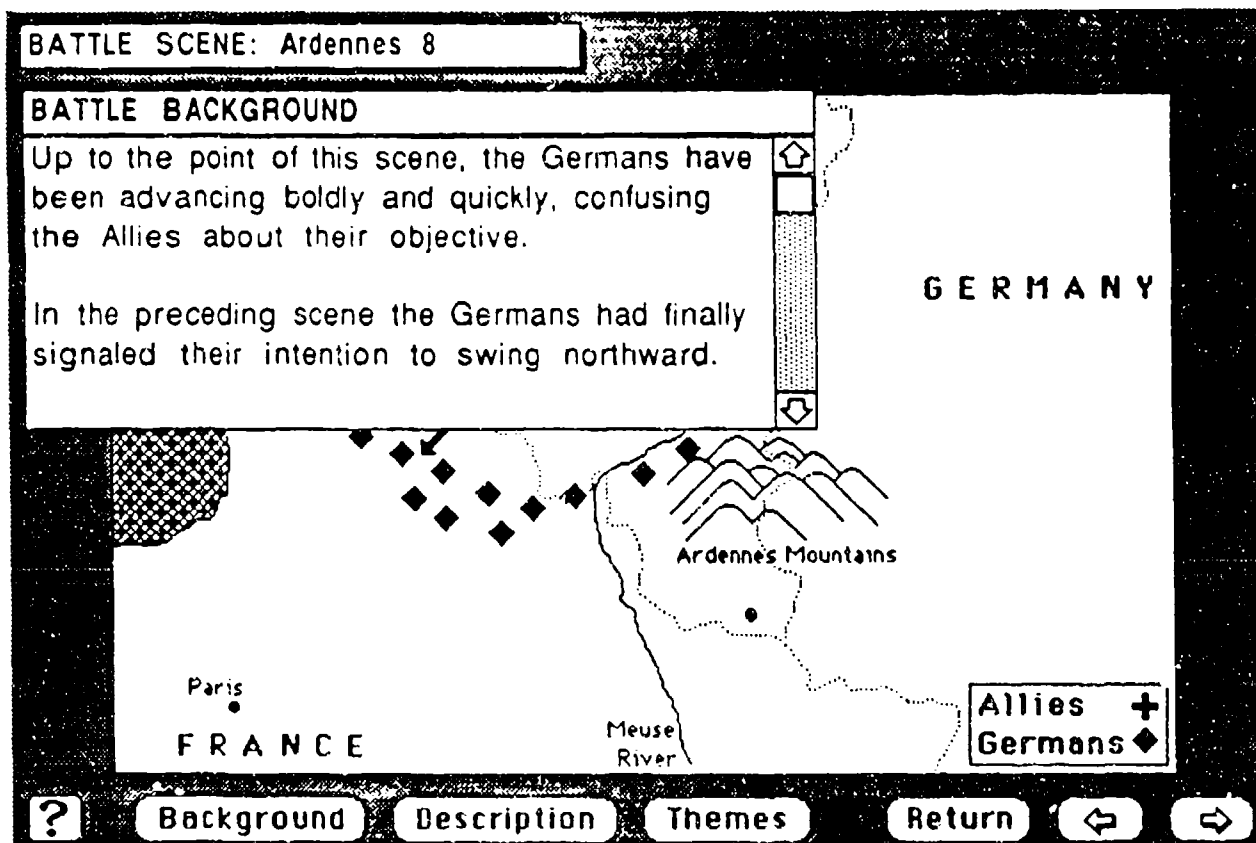


Figure 35. Background to Scene 8

If the option marked "Themes" is selected, a window is opened showing a list of only the SIA conceptual themes that are relevant to that particular scene. (See Figure 36 on the next page for an example.) Then, if the mouse is clicked on one of those themes, a *commentary* is provided (see Figure 37 on the *second* page below; only the first part of the commentary is provided--full commentaries will be illustrated later) that explains why that scene is considered to be an instance of that theme's application (i.e., an explanation of how the scene illustrates the SIA conceptual theme). This is a *crucial* aspect of Cognitive Flexibility Hypertexts because of the importance in Cognitive Flexibility Theory of teaching concepts *in use*. Rather than teaching each of the conceptual themes in isolation, they are always taught in an actual context of their application. Concept instruction is contextualized by having the description of the concept *tailored* to the particulars of the case scene, as well as by being presented in the context of other conceptual themes that co-occur with the concept in that same context (reducing the dangerous *compartmentalization* implicated in failures of AKA and that Cognitive Flexibility Theory reacts against).

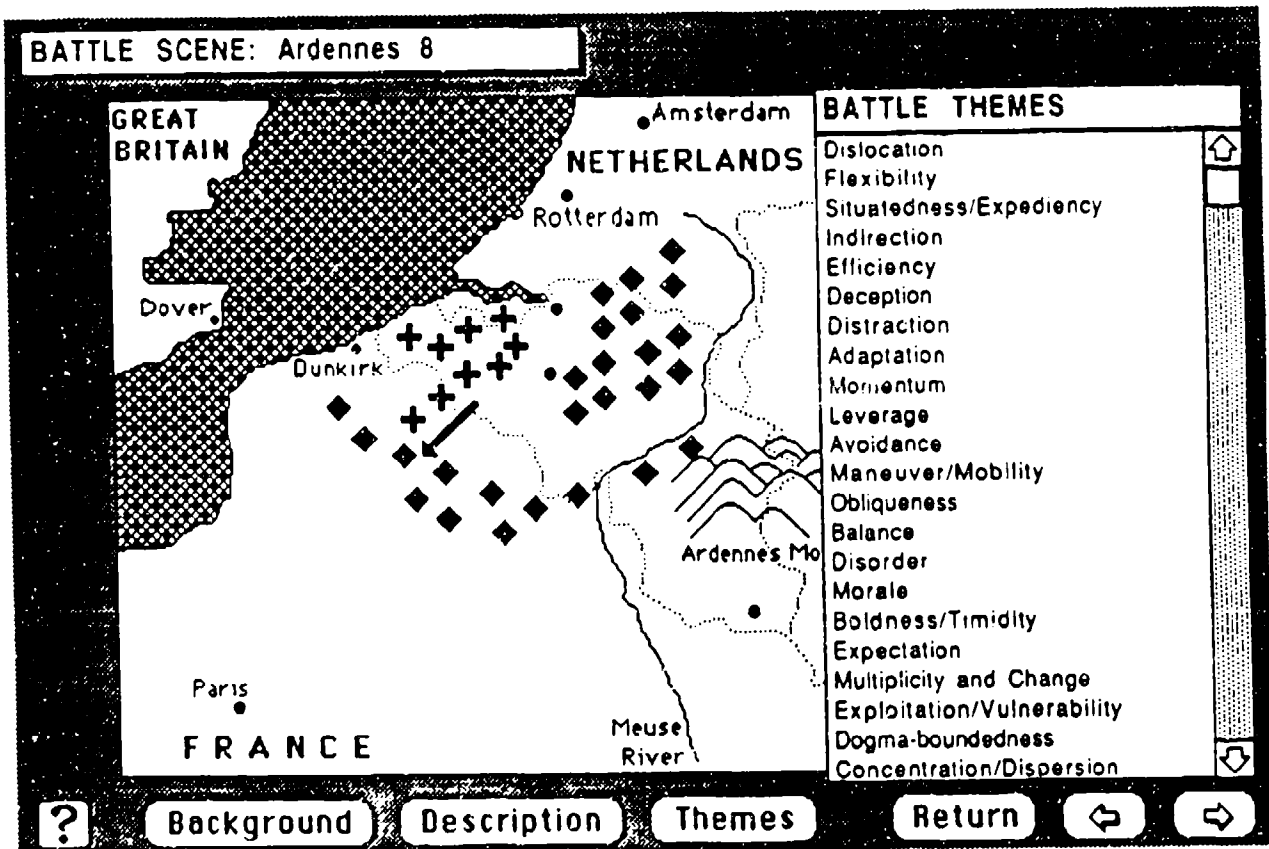


Figure 36. Themes for Scene 8

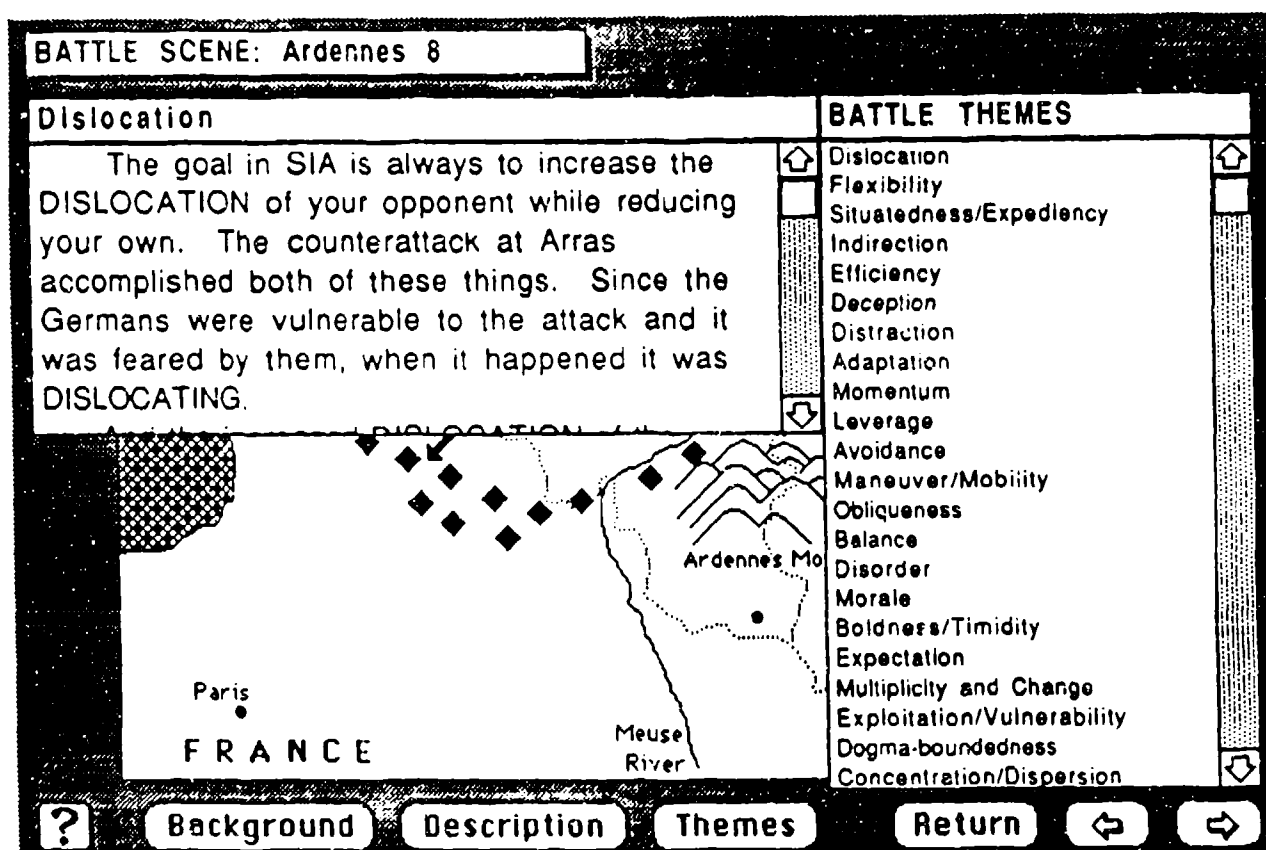


Figure 37. Example of Theme commentary

Across-scene options. A very important option is *theme-based search*. By selecting one or more themes from the list on the menu page (see Figure 38 on the next page), those scenes for which the selected conceptual themes are relevant are placed in a queue for sequential exploration by the learner. For example, the learner can choose to see just those scenes in Ardennes and Chancellorsville for which the *Flexibility* and *Boldness* themes of SIA are relevant. Note that it is possible to limit searches (as well as individual scene explorations) to any single case or subset of the complete set of cases. As we will see, the theme-search option is the basis for many theoretically important kinds of criss-crossing of the case information that makes up SIA's conceptual "landscape."

Click on the themes you want to search for. Then, click Initiate Search.

Core SIA Themes	Bridge Themes	
Dislocation	Deception	Balance
Flexibility	Mystification	Disorder
Situatedness/Expediency	Distraction	Morale
Indirection	Adaptation	Boldness/Timidity
Efficiency	Momentum	Expectation
	Leverage	Multiplicity and Change
	Avoidance	Exploitation/Vulnerability
	Maneuver/Mobility	Dogma-boundedness
	Obliqueness	

Initiate Search

Return to Main Menu

Case Themes	Case Features	
Concentration/Dispersion	Feints, Ruses, Traps	Central axis
Commitment	Shapes opponent	Breakthrough
Speed/Expeditiousness	Zig-zag	Penetration
Movement/Stasis	Counterattack	Wedging
Initiative	Threat: frontal	Rolling out
Security	Threat: flank	Force strength
Surprise	Threat: rear & communications	Terrain
Alternative objectives	Threat: envelopment	Weather

Figure 38. SIA theme-search menu

What are the SIA themes and how are they selected? Themes of wide scope for maximal transfer

SIA is an ill-structured domain, and therefore does not have a hierarchical structure with a single key conceptual entity superordinate to the rest. There are several ways one could fill in the sentence "SIA is a strategy of ____" (e.g., that it is a theory primarily of flexibility, or of deception, or of dislocation of the opponent). Since there are several themes of SIA that have considerable power in explaining battle cases as SIA events, but none is better than the others, *several* are chosen as *Core Themes* (see the upper left-hand corner of figure on the preceding page): 1) DISLOCATION of opponent (a psychological or physical imbalancing; a kind of "negative momentum") -- the positive side of this is the MOMENTUM of one's own; 2) FLEXIBILITY; 3) EXPEDIENT ADAPTATION TO SITUATIONS; 4) INDIRECTION (which includes MYSTIFICATION and SURPRISE); and EFFICIENCY. (The titles on the menu page are shorthands. The full titles are explained elsewhere in the program.)

Given the goal of widespread applicability of SIA knowledge across the range of cases that might be encountered (breadth of applicability or transfer), themes of SIA must be chosen which have *generality of application*. Since none of these themes is sufficient by itself, and *completeness* of case coverage is desired, there should be a small enough number of them so that they can *all* be kept in mind at all times. This requirement dictates that there should be a large number of them, but not more than can exceed the capacity of conscious attention. In summary, there should be as many as can be kept in mind at once (we use five Core Themes for SIA), they should be individually powerful in the breadth of their scope of application (i.e., cover a lot of cases), and they should overlap each other as little as possible (to maximize coverage for any single case). Which themes are selected is not crucial, so long as they roughly meet these conditions; they may not be a "perfect" set, but they will be a lot better than having a hierarchical organization with just one superordinate theme at the top (which is the most common organization of learning and instruction). No theme dominates the others. They are used in combination for *full analysis of both the individual case and SIA theory*.

Unfortunately, to achieve these conditions in an ill-structured domain you have to *sacrifice* something: *denotative specificity* of each theme. In order for a Core Theme to have the desired wide scope of application across cases (and flexibility of combination with other themes in analyzing an individual case), it must be accepted that instances of the themes will be diverse, united by only vague family resemblances (see the earlier discussions of conceptual variability in ill-structured domains). The Cognitive Flexibility Theory response to this problem, teaching with lots of facilitated experience with structural patterns of tailored use in cases, is taken up immediately below. [There are many other kinds of less generally applicable themes (but with greater denotative specificity) presented in the figure with the Theme Menu, but we mainly discuss the crucial Core Themes in this chapter; others are discussed briefly as opportunities arise, but are not systematically treated.]

How are SIA Themes taught? Case-based commentary plus theme search for conceptual variability.

If an SIA conceptual theme is not denotative, then what is its character and how can it be taught? As Cognitive Flexibility Theory indicates, an ill-structured conceptual entity's meaning

can not be determined without considering how it is *used* (see Wittgenstein's dictum that meaning is use, 1953). The structure of the Cognitive Flexibility Hypertext for SIA makes it very easy to teach the Core Themes in this way, by examining how they are used. For example, a first pass at teaching the theme of *Flexibility* (as it is used in the ill-structured theory of the SIA) would be to use the Theme Search option to identify those scenes in the Ardennes case that illustrate that theme, and to read their commentaries in succession. The commentaries presented on the next page show the outcome of that process -- what the Cognitive Flexibility Hypertext would show the learner given that request. In reading them, refer back to the maps and descriptions for each of the corresponding scenes.

Notice how the process of detecting patterns of family resemblance is aided by seeing several uses of the concept presented *in a row* (as contrasted with natural experience, which would not typically provide such convenient juxtapositions). Furthermore, each demonstration of the meaning of *Flexibility* in the SIA is *situated* within an actual *use* of the theme in a real battle scene, and it is accompanied by commentary that explains *how the concept is used in that context* (the Cognitive Flexibility Theory ideal of concepts-in-practice as the best preparation for the later use or transfer of the concepts).

Teaching Themes

It is time now to introduce the main conceptual themes of the Strategy of the Indirect Approach (SIA). Although more SIA themes will be introduced later, we begin here with just 4 themes that are especially central to SIA: 1) **DISLOCATION** of opponent (a psychological or physical imbalancing; a kind of "negative momentum") – the positive side of this is the **MOMENTUM** of one's own; 2) **FLEXIBILITY**; 3) **EXPEDIENT ADAPTATION TO SITUATIONS**; 4) **INDIRECTION/MYSTIFICATION/SURPRISE**; and **EFFICIENCY**.

One achieves the **DISLOCATION** of an adversary by using one's **FLEXIBILITY** to make an **EXPEDIENT ADAPTATION TO SITUATIONS** utilizing some kind of **INDIRECT** and/or **SURPRISING** course of action, preferably with the assistance of some already existing **DISLOCATION** of the opponent and **MOMENTUM** of one's own.

Rather than providing general definitions of these themes now that would not cover a wide enough range of examples, we will introduce each theme in the context of several actual examples of its occurrence. An important point that we want to make is that any SIA theme gets used in different ways at different times. An SIA theme does not have any single, general meaning. Consider as an example the theme of **FLEXIBILITY** as it is used in 3 different scenes from the battle of Ardennes.

BATTLE SCENE: Ardennes 3

BATTLE BACKGROUND

An anticipated German invasion of central Belgium has led to a massive rapid movement of Allied Forces to that region.

BATTLE DESCRIPTION

A German offensive force is reported in the Ardennes Mountains region (thought to be impassable by large armoured forces). The Allies ignore these reports and continue to stream with full commitment into central Belgium, north of the Ardennes.

Flexibility

Because the Allies were so certain that the Germans would attack through the central Belgian plain and they had focused all their thinking on that anticipated attack route, they were inflexible in the following way. They were unable to notice or credit reports of developments contrary to the anticipated ones. Inflexibility as seeing with the blinders of one's expectations on.

BATTLE SCENE: Ardennes 4

BATTLE BACKGROUND

The Allies have been responding en masse to an anticipated German attack in central Belgium. Reports of German forces in the Ardennes Mountains to the south have not been taken seriously.

BATTLE DESCRIPTION

It becomes clear that the Germans have launched a major offensive through the lightly guarded Ardennes region. This is particularly troubling because the overcommitted Allied forces are way out of position to react to the German move to the south (the most mobile Allied forces, having been sent into Belgium first, are the furthest in the wrong direction).

Also, the attack through the Ardennes simultaneously threatens: 1) the flank of the shaken Allied forces in Belgium, as well as their communications; 2) encirclement by a westerly followed by a northerly thrust; and 3) the vulnerable French capital of Paris.

Flexibility

The Allied reliance on one single plan in anticipation of German moves left them unable to respond quickly with new plans when it became clear that the Germans had violated their expectations. Inflexibility as rigid overreliance on just one plan or one way of thinking.

Also, by attacking through the Ardennes, the Germans had chosen a CENTRAL AXIS of attack from which multiple objectives could be targeted--flexibility as having viable alternatives availability to permit adaptive selection of the path of least resistance (AVOIDANCE).

BATTLE SCENE: Ardennes 5

BATTLE BACKGROUND

The Germans have broken through the lightly defended Ardennes in force, threatening multiple objectives (the flank of Allied forces in Belgium; encirclement of those forces; and vulnerable Paris).

The Allies are caught off balance from having thrust forcefully in the wrong direction and are out of position to counter the German move. They are shaken by the surprising and compromising turn of events.

BATTLE DESCRIPTION

The Germans move rapidly after breaking through at the Ardennes. They reach the Meuse River. Standard military procedure would be to wait for infantry and artillery to arrive and consolidate the position before crossing the river.

The German commander, placing a premium on the continuation of a rapid advance so as to further dislocate the Allies and to avoid momentum-halting contact, pushes immediately and successfully beyond the Meuse.

The Allies are still unable to respond. They are far out of position and off balance, uncertain where the Germans are going, and psychologically shaken.

Flexibility

The complete and immediate commitment of Allied forces to the central Belgian plane because of the strong expectation that the Germans would attack there reduced the physical flexibility of the Allies if they turned out to be wrong about German intentions (which they were). In inverse proportion to the strength of their response in central Belgium, they were weakened in the south by their inability to quickly redirect their forces (see JIU-JITSU EFFECT). Inflexibility is seen as the physical limitations on the adaptive redirection and relocation of forces (e.g., momentum in wrong direction; being off-balance).

BATTLE SCENE: Ardennes 6

BATTLE BACKGROUND

The Germans have been advancing rapidly and virtually unopposed following the Ardennes breakthrough. The Allies have been unable to respond.

BATTLE DESCRIPTION

The German army disperses somewhat and begins to follow a zig-zag, unpredictable course.

The Allies continue to be unable to make a credible response.

Flexibility

The dispersed, zig-zagging advance of the Germans reflected flexibility in several ways. First is the physical ability of forces to follow a rapid zig-zagging course--flexibility as physical mobility. Second, by confusing the Allies about their direction of advance and thereby paralyzing them, the Germans were able to AVOID confrontations with blocking forces--flexibility as freedom to move without hindrance from the opponent. Dispersion promotes flexibility by enabling the option of fluid re-concentration in a variety of directions and by decentralizing command.

BATTLE SCENE: Ardennes 8

BATTLE BACKGROUND

Up to the point of this scene, the Germans have been advancing boldly and quickly, confusing the Allies about their objective.

In the preceding scene the Germans had finally signaled their intention to swing northward.

BATTLE DESCRIPTION

The Allies launch a small counterattack at Arras. The Germans, not knowing the size of the attack and fearing their exposed flanks, halt their forward advance.

Flexibility

SIA is a theory of FLEXIBILITY. You must always retain FLEXIBILITY--it is a precondition of SIA action. And you must respond FLEXIBLY to circumstances.

In this scene, the Allies were FLEXIBLE. They were able to promptly respond to the noticed change in the SITUATION with an EXPEDIENT (SITUATIONALLY appropriate) change of their own. FLEXIBILITY as the ability to change what you've been doing in response to changing circumstances.

The Allies had increased FLEXIBILITY because the Germans' slowdown and concentration in Scene-7 made it clear that the Allies did not have to worry as much about defense in other sectors. FLEXIBILITY as the ability to free up scarce resources for new initiatives.

The counterattack reduced the Germans' FLEXIBILITY. Because they worried about their vulnerability at that point, and they lacked information about the magnitude of the attack, they were forced to respond to it rather than pursuing other initiatives. INFLEXIBILITY as compelling a course of action from an opponent.

[See ADAPTATION, CHANGE, LEVERAGE, MANEUVER, AVOIDANCE, VULNERABILITY]

The same process would be followed to *begin* to teach the other Core Themes. The word "begin" is italicized to highlight an important point about Cognitive Flexibility Hypertext instruction: it begins with microcosms and then expands outward to larger units of study. The "shape" of instruction in a Cognitive Flexibility Hypertext is a kind of "spiralling outward" from small detailed treatments of segments of individual cases. So, in using the SIA program, the learner begins with just the Ardennes campaign, which is first covered superficially, and then broken down into segments or scenes that are covered in increasingly greater detail, first individually, and then in relation to each other. For example, the learner seeing Scene 7 would cover just four or five SIA themes on the first pass, with subsequent readings peeling of increasing layers of complexity. Only after using this one campaign as a site for complexity instruction for a large amount of time would other battles be introduced. [In order to allow for large numbers of instructive case-comparisons (criss-crossings), subsequently introduced battles are given much more superficial coverage in the SIA Cognitive Flexibility Hypertext than Ardennes does.] So, the themes would be revisited in a similar manner later, but crossing different battle cases rather than just crossing the scenes of one battle case (Ardennes). Similarly, other processes of Cognitive Flexibility Theory-based use of the SIA hypertext for the Ardennes scenes would also be recapitulated with a larger "grain size" at a later stage of knowledge acquisition (when the learner is better prepared to deal with larger "bytes" of information).

How are cases and scenes taught? The Central Role of the Minicase in Introducing Complexity Early in Instruction

One of the most important lessons from the body of empirical findings in the part of our research program concerned with what goes wrong in AKA was that concerning the failure of traditional incrementalism (i.e., the movement from simple to more complex treatments of conceptual knowledge). When you "start simple", characteristics of those early, simplified learnings *block* the later attainment of complexity (see Section 3). What is the alternative suggested by Cognitive Flexibility Theory? Begin with case treatments that retain the important characteristics of complexity (e.g., multiple interacting conceptual themes). But how can this be done without overwhelming the learner? By beginning with very small segments of cases, what we call "bite-size chunks" of complexity or "minicases." By reducing the demands on the learner to keep track of case details, more processing capacity remains available to process aspects on conceptual complexity. We refer to this alternative approach to sequencing, moving from smaller to larger staging grounds for complexity rather than from oversimplified to more complex treatments, as a "New Incrementalism" (see Chapter 4.4).

We will illustrate with Scene 8 from the Ardennes case. Consider how much is learned just by reading the Core Themes commentaries for this scene. [The reader should definitely re-read the Background and Description information for this scene again before reading the following discussion of this late scene in the Ardennes campaign. They are reproduced on the next page. The battle map for the scene is reproduced on the following page, Figure 39. The Core Theme commentaries, as they would be presented if clicked upon, are on the pages following the scene map.]

BATTLE SCENE: Ardennes 8

BACKGROUND

The Germans have continued to perplex the Allies by pursuing an erratic direction of advance between a few alternative objectives. The Allies have been unable to respond.

SCENE DESCRIPTION

The Germans consolidate their southern army, reducing its dispersion and zig-zagging, and finally signaling their objective: to swing north and entrap the Allied forces between the southern army and the army that had attacked first, to the north, and which is advancing westward through the low countries (the central Belgian plain and the Netherlands).

The Allies launch a small counterattack at Arras. The Germans, not knowing the size of the counterattack and fearing their exposed flanks, halt their forward advance.

As the Germans pause to consolidate, the Allies are given important time to recover their balance and plan for the extrication of their trapped forces.

The successful evacuation of the trapped Allied forces at Dunkirk followed.

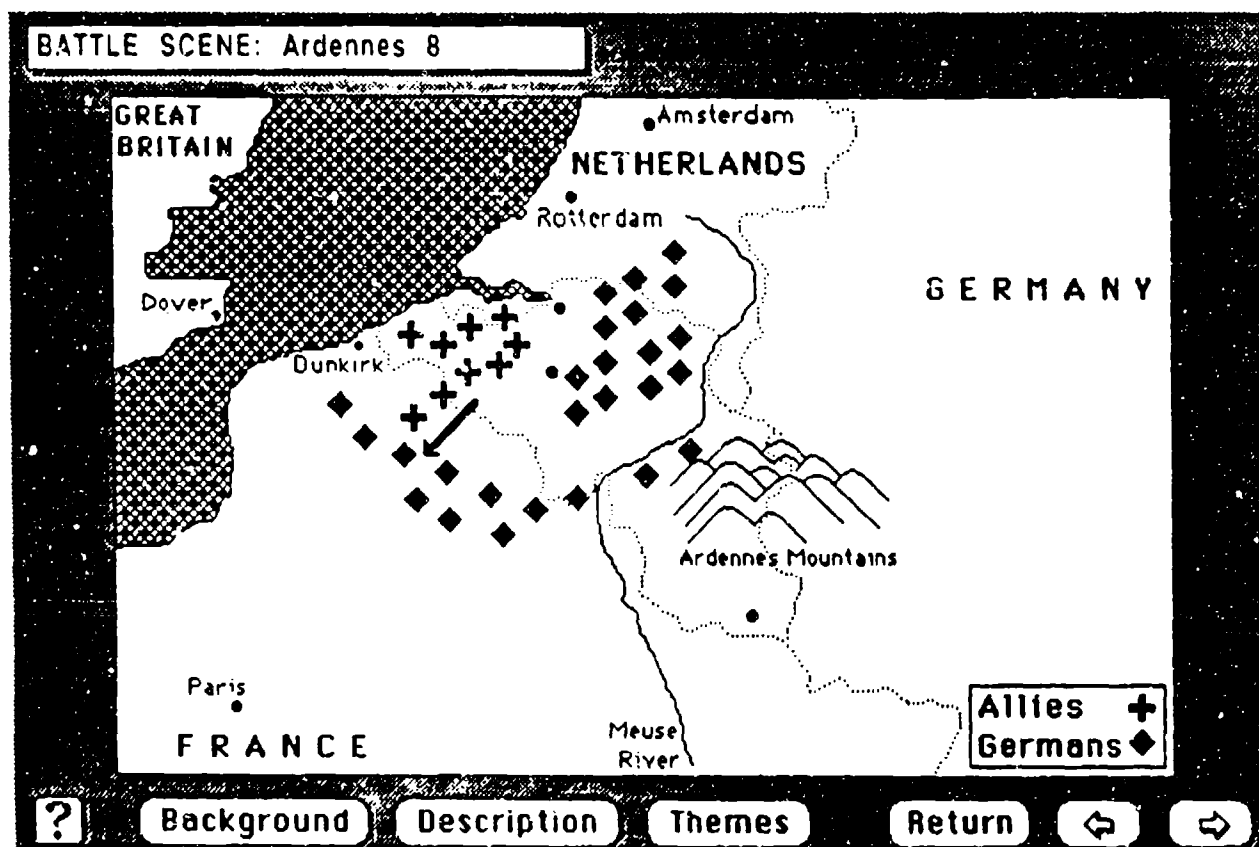


Figure 39. Map for Scene 8

Continuing our closer look at the operation of the themes of the Strategy of the Indirect Approach, following the demonstration of the complexity of themes in SIA (Flexibility), we will examine the complexity of cases as they relate to SIA. Scene-8 will be examined in depth to show how several SIA themes get applied to any particular case, how the SIA themes (including some new ones) fit into (are tailored to) an actual case, how they are interrelated, and so on. In other words, not only is a detailed case study presented, but a demonstration is provided of how one needs to think about battles in order to apply SIA. [It should be noted that Scene-8 is the first instance during the Ardennes campaign of the Germans violating SIA and the Allies adhering to it!]

We begin with the application of some of the SIA *Core Themes* to Scene 8. (The in-depth analysis of this scene will be continued in subsequent sections.)

----- Situatdness/Expediency

SIA is a theory of close analysis of SITUATIONS and EXPEDIENT tailoring of responses to them. In Scene-7, the SITUATION CHANGED. The Germans halted their drive/penetration in order to concentrate their forces to protect vulnerabilities. The Allies noticed and reacted to this opportunity. SITUATED EXPEDIENCY as adaptation in the context of attention to change in the situational condition (made possible by FLEXIBILITY).

[See MOMENTUM, LEVERAGE, VULNERABILITY, ADAPTATION, CHANGE]

--- Flexibility

SIA is a theory of FLEXIBILITY. You must always retain FLEXIBILITY--it is a precondition of SIA action. And you must respond FLEXIBLY to circumstances.

In this scene, the Allies were FLEXIBLE. They were able to promptly respond to the noticed change in the SITUATION with an EXPEDIENT (SITUATIONALLY appropriate) change of their own. FLEXIBILITY as the ability to change what you've been doing in response to changing circumstances.

The Allies had increased FLEXIBILITY because the Germans' slowdown and concentration in Scene-7 made it clear that the Allies did not have to worry as much about defense in other sectors. FLEXIBILITY as the ability to free up scarce resources for new initiatives.

The counterattack reduced the Germans' FLEXIBILITY. Because they worried about their vulnerability at that point, and they lacked information about the magnitude of the attack, they were forced to respond to it rather than pursuing other initiatives. INFLEXIBILITY as compelling a course of action from an opponent.

[See ADAPTATION, CHANGE, LEVERAGE, MANEUVER, AVOIDANCE, VULNERABILITY]

--- Indirection

SIA theory prescribes that you attain your goals through INDIRECT actions. A flank counterattack is by definition, literally, an INDIRECT move. It is nonfrontal. The Allies were in no condition to meet the Germans frontally. However, an INDIRECT attack on their exposed flank required fewer forces

(increasing the EFFICIENCY). By having ceased their indirect movement by slowly and deliberately directing an approach to one objective, the Germans' directness made AVOIDANCE impossible.

Less literally, the counterattack was DIRECT in that it was expected. This DIRECT aspect is more than balanced by the German vulnerability to the move and the great potential it therefore had for alteration in the DISLOCATION quotient for the two forces.

[See LEVERAGE, OBLIQUENESS, VULNERABILITY]

Dislocation

The goal in SIA is always to increase the DISLOCATION of your opponent while reducing your own. The counterattack at Arras accomplished both of these things. Since the Germans were vulnerable to the attack and it was feared by them, when it happened it was DISLOCATING.

And the increased DISLOCATION of the Germans that the counterattack produced kept them from resuming their initiative, thus allowing the Allies more time to recover (i.e., it reduced further the Allies' own highly compounded DISLOCATION). This eventually contributed to the Allies' ability to evacuate many of their troops at Dunkirk.

[See VULNERABILITY, MYSTIFICATION, LEVERAGE, BALANCE, BOLDNESS/TIMIDITY]

The Germans had conducted a series of MANEUVERS that had successively confused (MYSTIFICATION) the Allies, throwing them off BALANCE and reducing their MORALE. Each step in this series compounded the DISLOCATION of the Allies. When an opponent is already reeling, further shocks have an enhanced effect. This compounded DISLOCATION was seriously lessened, however, when the Germans relinquished some of their MOMENTUM by slowing their advance in Scene-7. The reduced MOMENTUM allowed the Allies an opportunity to counterattack a VULNERABLE German flank. Because the counterattack was so feared by the Germans (TIMIDITY), it had the effect of nearly halting the German MOMENTUM.

Now look again at the Background, Description, and map for Scene 8. Are you not beginning to develop a deeper understanding of this example and of how the Strategy of the Indirect Approach works? [To reinforce this lesson one more time, you may want to quickly review the thematic analysis of Scene 8 that was just presented.]

Beginning Instruction with Features of Complexity Represented from the Outset: An Illustration of the "New Incrementalism" of Instructional Sequencing. Notice how the complexity of SIA has begun to be illustrated without overwhelming the learner. Using just a small amount of case information as background (the few lines of the description of Scene 8), a lesson has already been taught about the necessity of applying multiple perspectives to a case. The learner sees from the beginning that an example is not an example of just one major conceptual point. Scene 8 has been *seen* as an instance of SIA in four partly overlapping but different ways. Not only do the multiple SIA perspectives show how rich SIA is (SIA is not any *one* dominant thing, but the interplay of several), but an important point has been made about how much there is to see in a case, and how much would be missed in fully covering a case by trying to describe it in just *one* way. And a lesson about each of 4 of the Core Themes was taught. Another lesson was taught about the tailoring of abstract themes to their contexts of case application (concepts-in-practice). The presence of cross-references in the theme commentaries began to teach the Cognitive Flexibility Theory tenet of non-compartmentalization, and to provide important clues to the learner about where to explore next.

All of these lessons, embodying different *general* aspects of complexity and Cognitive Flexibility Theory, as well as many *specific* lessons about this one concrete scene, were conveyed in a *short amount of time*, and with *minimal demands on the learner* (at least with respect to the amount of case information that must be kept in mind--again, because the concrete case context for the lessons is so slender and easy to master). That the demands are not too great is especially so considering that these Core Theme commentaries will be re-read later, in other contexts, so it is not necessary for the learner to fully learn and remember what was learned about the themes--more opportunities will present themselves to deepen their understanding. This spiralling depth of coverage on subsequent re-visitations is an important feature of Cognitive Flexibility Theory instruction. We have referred to the gradual build-up of complexity from initial presentations that have the features of complexity but staged in a manageable context as the "New Incrementalism" of instructional sequencing -- *from the very beginning*, the learner sees that learning the domain involves such Cognitive Flexibility Theory principles of multiple conceptual representations, interconnectedness instead of compartmentalization, concepts tailored to context rather than by the application of some definition in the abstract, and so on. In traditional "Simple to Complex" instructional sequencing, the opposite, more well-structured pole of these dimensions would characterize initial instruction; and we have seen repeatedly demonstrated (see Section 3), when you start simple it often blocks the later acquisition of complexity, and can even instill epistemological beliefs and habits of learning that are *antithetical* to learning complexity and ill-structuredness. As this section continues, the way that additional complexity is incrementally added from this "bite-size chunk" will become clear (e.g., more theme commentaries will be added to those of the Core Themes, and more interrelational weaving will occur, still with the same minicase as the staging ground).

Effects of Adding Minicases: Compacted and elaborated lessons about SIA from tracking just 3 themes through several SIA scenes.

We have seen a surprising amount of case-specific and more general Cognitive Flexibility Theory lessons taught using just one small scene (minicase) from the larger Ardennes case. Now see what happens when you do the same thing for a few other small scenes from Ardennes. Read through the Core Theme commentaries (all but *Efficiency* are included) on the pages below for the rest of the Ardennes campaign. For purposes of exposition, the 8 scenes have been consolidated to just 4 (corresponding to the several Ardennes scenes pictured and described initially, which should be reviewed before reading the material below). The proliferation of lessons should be obvious.

Consolidated Narrative of Ardennes Case in Terms of Some of the SIA Core Themes

BATTLE BACKGROUND

A sizable German force attacked through central Belgium, as the Allies had strongly anticipated. Allied forces immediately streamed northward en masse towards the area of the attack in massive numbers, with their most mobile forces (the ones that could get their quickest) leading the way.

BATTLE DESCRIPTION

A German offensive force is reported in the Ardennes Mountains region (thought to be impassable by large armoured forces), to the south of central Belgium. The Allies do not take these reports seriously, and continue to stream with full commitment into central Belgium, north of the Ardennes.

[FLEXIBILITY: The Allies' rigid viewing of the unfolding events in terms of their expectations prevented them from processing things that deviated from that scheme (like the German movement in the Ardennes, which didn't fit the Allied plan at all).]

[EXPEDIENT ADAPTATION: The Allies will not be able to FLEXIBLY and ADAPTIVELY react to changing circumstances if circumstances are not being monitored (especially for changes from expectation). SIA and the history of battle says that one must always be ready to adapt to changing and unexpected circumstances -- the INFLEXIBLE and limiting viewpoint of the Allies violated this tenet.]

[INDIRECTION/MYSTIFICATION/SURPRISE: The Allies assumed a direct, frontal posture towards the known German force, expecting a direct, frontal attack. This left them unprepared for and surprised by the Germans' unexpected Indirect movement through the Ardennes.]

Later, it becomes clear that the Germans have launched a major offensive through the lightly guarded Ardennes region.

[INDIRECTION/MYSTIFICATION/SURPRISE: The Germans chose a difficult attack route through mountains instead of the direct path of least effort through the central Belgian plain that had been expected.]

[EXPEDIENT ADAPTATION: The Germans attacked where the Allied line of defense was weak.]

This is particularly troubling for the Allies because the overcommitted Allied forces are way out of position to react to the German move to the south (the most mobile Allied forces, having been sent into Belgium first, are the furthest in the wrong direction).

[FLEXIBILITY: The Allies have seriously reduced their flexibility of response by so fully committing to the expected northern attack. They have no plans for other contingencies (like the one that occurred).]

[DISLOCATION/MOMENTUM: The fact that the Allies guessed so totally wrong, committed so much based on that guess, and have no immediate capability of adjustment is very upsetting, throwing the commanders into a situation of confusion and turmoil as they hurriedly try to come up with a new plan to meet the unexpected German thrust. The Allies being so far out of position, so imbalanced, is a physical Dislocation. Surprise plus being out of position is Dislocating. The Allied commanders are shaken!]

[EXPEDIENT ADAPTATION: Although the changed situation is now recognized by the Allies, their INFLEXIBILITY has produced a physical dislocation of their forces leading to an inability to ADAPT to those changes. Also, the unsettling DISLOCATION and lack of an alternative contingency plan will make it harder to adapt psychologically (strategically) to the changed situation.]

The attack through the Ardennes simultaneously threatens 3 potential German objectives: 1) the vulnerable flank (side) of the shaken Allied forces in Belgium, as well as their lines of communication and command back to France;

2) encirclement of the Allied forces in Belgium by a westerly movement followed by a northerly thrust; and 3) the vulnerable French capital of Paris.

[FLEXIBILITY: The Germans have several options as to where they move next.]

[DISLOCATION/MOMENTUM: The psychological disturbance of the Allies produced by the unexpected and unprepared-for Ardennes offensive is compounded by the importance of the three alternative objectives the Germans threaten and the uncertainty of which objective they will choose (making planning for a response difficult) -- important vulnerabilities and uncertainties about the enemy's course of action both increase Dislocation.]

[ADAPTIVE EXPEDIENCY: The Germans can choose an option depending on what the Allies do. For example, if the Allies move to protect their flank, the Germans could then move on Paris. Whatever objectives the Allies guard against, the Germans can then adaptively move against one of the other objectives. The Allies, on the other hand, must adapt to whatever decision the Germans make, while the Germans get to pick whatever's best for them -- there is a greater burden of ADAPTATION on the Allies.]

BATTLE BACKGROUND

The Germans have broken through the lightly defended Ardennes in force, threatening multiple objectives (the flank of Allied forces in Belgium; encirclement of those forces; and vulnerable Paris).

The Allies are caught off balance from having thrust so forcefully in the wrong direction and are out of position to counter the German move. They are shaken by the surprising and compromising turn of events.

BATTLE DESCRIPTION

The Germans move rapidly after breaking through at the Ardennes. They reach the Meuse River. Standard military procedure would be to wait for infantry and artillery to arrive and consolidate the position before crossing the river. The German commander, placing a premium on the continuation of a rapid advance so as to further dislocate the Allies and to avoid momentum-halting contact, pushes immediately and successfully beyond the Meuse. The Allies are still unable to respond. They are far out of position and off balance, uncertain where the Germans are going, and psychologically shaken.

[DISLOCATION/MOMENTUM: The Germans capitalize on their Momentum (and the Allies' corresponding Dislocation) by pursuing their advantage very boldly.

The German army then disperses somewhat (spreads out) and begins to follow a zig-zag, unpredictable course. The Allies continue to be unable to make a credible response.

[FLEXIBILITY: A force that is Dispersed can be reoriented in more directions than a concentrated force.]

[DISLOCATION/MOMENTUM: Allied DISLOCATION, already high, is increased by the additional confusion and uncertainty caused by the unpredictable course taken by the Germans (rapid, zig-zagging, and, because of dispersion, capable of being sent in a variety of directions). German MOMENTUM is protected by the various things they are doing to Avoid an Allied counter-attack (an INDIRECT approach): dispersed forces offer no concentrated target; rapid movement makes for a difficult, moving target; and zig-zag course makes for unpredictability of where enemy will be next.]

[ADAPTIVE EXPEDIENCY: ADAPTATION is a problem: you can't adapt to what you can't see; the Allies can't adapt if they don't know what they are adapting to.]

[INDIRECTION/MYSTIFICATION/SURPRISE: In order to maintain their MOMENTUM, the Germans do various things to AVOID the Allies -- avoidance as a strategy (here, to put the Allies in an even worse predicament later on) is the ultimate form of INDIRECTION. Note also that the AVOIDANCE is facilitated by the MYSTIFICATION produced by the zig-zagging, unorthodox rapid advance without consolidating positions, and dispersion of forces.

BATTLE BACKGROUND

The Germans have continued to perplex the Allies by pursuing an erratic direction of advance between a few alternative objectives. The Allies have been unable to respond.

BATTLE DESCRIPTION

The Germans consolidate their southern army, reducing its dispersion and zig-zagging, and finally signaling their objective: to swing north and entrap the Allied forces between the southern army and the army that had attacked first, to the north, and which is advancing westward through the low countries (the central Belgian plain and the Netherlands).

The Allies launch a small counterattack at Arras. The Germans, not knowing the size of the counterattack and fearing their exposed flanks, halt their forward advance.

[INDIRECTION/MYSTIFICATION/SURPRISE: Any attack to the flank (side) of a force, rather than to its front, is Indirect. SIA theory prescribes that you attain your goals through INDIRECT actions. A flank counterattack is by definition, literally, an INDIRECT move. It is nonfrontal. The Allies were in no condition to meet the Germans frontally. However, an INDIRECT attack on their exposed flank required fewer forces (increasing the LEVERAGE/EFFICIENCY of the counter-attack -- the ability of that counter-attack to have a greater effect on the Germans than it should have, given the size of the attack). The counterattack was SURPRISING because it followed a long pattern of the Allies taking no substantial action at all to disturb the Germans' progress.]

[ADAPTIVE EXPEDIENCY: SIA is a theory of close analysis of SITUATIONS and EXPEDIENT tailoring of responses to them. In Scene-7, the SITUATION CHANGED. The Germans halted their drive/penetration in order to concentrate their forces to protect vulnerabilities. The Allies noticed and reacted to this opportunity. SITUATED EXPEDIENCY is used here as adaptation in the context of attention to change in the situational condition (made possible by FLEXIBILITY).

[FLEXIBILITY: In this scene, the Allies were FLEXIBLE. They were able to promptly respond to the noticed change in the SITUATION with an EXPEDIENT (SITUATIONALLY appropriate) change of their own. FLEXIBILITY as the ability to change what you've been doing in response to changing circumstances. The Allies had increased FLEXIBILITY because the Germans' slowdown and concentration in Scene-7 made it clear that the Allies did not have to worry as much about defense in other sectors. FLEXIBILITY as the ability to free up scarce resources for new initiatives.]

[DISLOCATION/MOMENTUM: The goal in SIA is always to increase the DISLOCATION of your opponent while reducing your own. The counterattack at Arras accomplished both of these things. Since the Germans were vulnerable to the attack and it was feared by them, when it happened it was DISLOCATING. And the increased DISLOCATION of the Germans that the counterattack produced kept them from resuming their initiative MOMENTUM, thus allowing the Allies more time to recover (i.e., it reduced further the Allies' own highly compounded DISLOCATION). This eventually contributed to the Allies' ability to evacuate many of their troops at Dunkirk. A force with MOMENTUM operating against a DISLOCATED enemy should keep the pressure on; relinquishing that MOMENTUM to consolidate gains out of TIMIDITY and because of an overreaction to the counter-attack was a grave error.

BATTLE BACKGROUND

The Allies have finally broken the German momentum, capitalizing on knowing where the German army was heading and the fact that the Germans had ceased their rapid movement to launch a small but very effective counterattack. In response, the Germans have halted their advance to consolidate their position and defend their flank.

BATTLE DESCRIPTION

As the Germans pause to consolidate, the Allies are given important time to recover their balance and plan for the extrication of their trapped forces.

(DISLOCATION/MOMENTUM: The Germans' sacrificing of their Momentum in a timid pursuit of security allowed the Allies to recover from their Dislocation.)

The successful evacuation of the trapped Allied forces at Dunkirk followed.

You should be experiencing a deepening awareness that your understanding is getting richer, more intimate, more complex, just from this very *limited* exploitation of the features of the Cognitive Flexibility Hypertext useful for exploring that complexity of understanding that advanced knowledge acquisition increasingly requires. And, again, the scenes of the Ardennes campaign are a fairly compact staging ground for teaching elaborated lessons about SIA. There is not that much factual detail that must be kept in mind while tracking these illustrations of SIA themes and their interplay. Yet, consider how much about SIA and about its application to this particular case (the Ardennes campaign) is learned from just this small amount of easily managed case information, in relatively short amounts of time (keeping in mind, again, that not all of this would be covered in a row, all at one time; however, the amount of time it would take across occasions is very small). And, as we will see below, there are many more lessons that could potentially be taught, as the design features of Cognitive Flexibility Hypertexts are exploited to teach other Cognitive Flexibility Theory lessons using the *same set of minicases*.

Going Beyond the Core Themes: The Effects of Adding Themes on Successive Passes

Now that we have seen the powerful effects on increasing the *fullness of scene coverage* and producing *additional interconnection/assembly paths* for future use of having just a few Core Themes used to analyze the SIA aspects of a battle scene, we will now show how much gets added to instruction, with just the same short minicase-scene, as additional theme commentaries are explored on later visits to the scene. [Scene 8 will again be explored in-depth, in various new ways. That scene's earlier exploration of Core Theme commentaries, as well as its map and description, should be reviewed again at this time.]

Now let's look in even greater depth at how the Strategy of the Indirect Approach is applied to Scene 8. We will look at the commentaries for some additional, non-core themes (called bridge themes because they help to connect the more abstract Core Themes to specific, denotative case details, and they connect Core Themes with each other). During this further exploration of Scene 8 we will also occasionally *return* to Core Theme commentaries we have already read because of the new light they cast in the context of being juxtaposed with different themes than they were earlier.

Leverage

The Allies were able to seriously disrupt the German advance with just a small counterattacking force. The DISLOCATION for the Germans caused by the counterattack was out of proportion to the risk that it posed--this is an example of LEVERAGE.

This LEVERAGING was made possible by several things:

First, the Germans had, in the preceding Scene-7, slowed the SPEED of what had earlier been a continuously rapid advance-- a reduction in MOVEMENT and INITIATIVE that was a CHANGE from the most recent pattern. The fact that it was a CHANGE increased its effects. The reduction in INITIATIVE allowed the Allies to recover from some of their recent DISLOCATION. The reduction in MOVEMENT provided a VULNERABLE target to hit at. The moral of this story is that an attacking force that has MOMENTUM should not let its opponent "off the hook."

Second, the Allies used RUSES [L4] to DECEIVE [L2] the Germans about the size of their counterattacking force. Of course, the (misperceived) larger size of the force caused the Germans' FEAR to be increased and therefore increased the LEVERAGE (increased the DISLOCATION).

Third, the effect of the small counterattacking force was LEVERAGED by being a nonfrontal (OBLIQUE) attack. More forces are always required for a frontal attack. By attacking INDIRECTLY (towards a flank), less force was required to have a big effect.

Fourth, A PENETRATING column is VULNERABLE to FLANK COUNTERATTACKS. By attacking a VULNERABLE point the Allies were able to LEVERAGE their attack (i.e., have a greater effect than usual for a small force). This was an EFFICIENT way to EXPLOIT the opportunity presented by the Germans slowing their advance. They had slowed their advance because the textbook VULNERABILITY of a PENETRATING column to a COUNTERATTACK TO FLANK made them afraid (TIMIDITY). CONCENTRATING is the recommended solution in such a case, and it was incompatible with further MOVEMENT and INITIATIVE, thus reducing the Germans' MOMENTUM.

Fifth, even the small amount of forces made available for the counterattack would not have been possible if the Germans hadn't allowed the Allies to temporarily relax elsewhere. The halt in the German INITIATIVE gave the Allies more FLEXIBILITY to redirect some of their retreating and defending forces towards an offensive action (LEVERAGING in one place by relieving need in another place.)

Boldness/Timidity

The Germans were afraid that the Allies would attack their VULNERABLE flank. As a result, they stopped their BOLD advance, which had been so successful, and protected against their fears. This TIMID action was a kind of

self-fulfilling prophecy: by lessening their INITIATIVE they enabled the Allies to actually make the counterattack that they feared (w/ max. LEVERAGE, given the combination of VULNERABILITY + TIMIDITY + STATIONARY). When one has DISLOCATED an enemy by BOLD actions, it's a mistake to become TIMID and allow the opponent to recover from their DISLOCATION. If you've got them down, don't let them get back up again.

Deception

The Allies used RUSES to DECEIVE the Germans about the size of their counterattacking force. The fact that the Germans had EXPECTED such a counterattack made them more susceptible to the DECEPTION. The DECEPTION increased the DISLOCATION produced by the small counterattacking force, making the action more EFFICIENT (a LEVERAGING effect).

Avoidance

The counterattack thwarted a major goal of the Germans: AVOIDING confrontations with the Allies in order to continue their PENETRATION. By MANEUVERING to AVOID battle, the Germans had been able to stay one step ahead of the Allies, MYSTIFYING them about German intentions (ALTERNATIVE OBJECTIVES), throwing them off BALANCE, and causing DISORDER, all of which had cumulatively DISLOCATED the Allies. These effects were seriously diminished when the Germans slowed down in Scene-7 allowing the Allies an opportunity to attack them. Much of the MOMENTUM that the Germans had built was lost, eventually contributing to the Allies successful evacuation at Dunkirk.

The Germans' SPEED + MOBILITY + DISPERSION + UNPREDICTABILITY had enabled AVOIDANCE. In Scene-7 they had eliminated all of these: they slowed down, reduced movement, concentrated, and became more predictable. As a result, it was easier for the Allies to force non-AVOIDANCE.

The Allies were smart to force the Germans not to AVOID them because that was what the Germans wanted to do--the Germans were the INITIATOR, they had the MOMENTUM. It is usually a good idea to try to thwart the desires of the INITIATOR or MOMENTUM-HOLDER.

The decision not to allow the Germans to AVOID the Allies was good also because it was something that SCARED the Germans, something they thought they were especially VULNERABLE to, and thus was very DISLOCATING to them.

It was more acceptable for the Allies to not AVOID the Germans to the extent that their MANEUVER was OBLIQUELY CONCENTRATED against a WEAKER FORCE (EFFICIENCY) and FEARED. These are also conditions of EXPEDIENCY!!

The counterattack engagement (non-AVOID) SHAPED THE OPPONENT (SECURITY). The Allies knew where the Germans were (leading to greater FLEXIBILITY & EFFICIENCY for the Allies--they didn't have to distribute everywhere to protect against everything. And it forced the Germans to be where the attack was occurring, reducing their FLEXIBILITY.

AVOIDANCE is related to many of the Core Themes: Avoidance is INDIRECTION (the German offensive was thus inherently INDIRECT). To achieve Avoidance, you need FLEXIBILITY (the slowdown for concentration in Scene-7 reduced the Germans' FLEXIBILITY). To achieve Avoidance, you also need to pay attention to all of the things going on at any one time and to the rapid change in what's going on (SITUATEDNESS). If it's already known that the opponent wants to Avoid you, and you hit him, that's DISLOCATING.

Maneuver/Mobility

In this scene the Germans paid a price for stopping their high tempo of MANEUVER in the preceding Scene-7. The Germans owed much of their earlier success to their highly mobile MANEUVERING; they followed a ZIG-ZAG course, they CONFUSED the Allies about their intentions by first threatening one way and then another (ALTERNATIVE OBJECTIVES), they used RUSES to DECEIVE the Allies about their position (SECURITY), and so on. The MANEUVERING was FLEXIBLE and SITUATIONALLY EXPEDIENT (ADAPTIVELY altering direction as circumstances dictated). It built MOMENTUM and resulted in compounded DISLOCATION. The Allies simply became paralyzed after having been thrown off BALANCE so often. The success of the German MANEUVERS decreased the ability of the Allies to MANEUVER in response. By changing strategy from FLEXIBLE MANEUVER to a directed CONCENTRATION in Scene-7, they were losing much of the INDIRECTNESS that had allowed them to AVOID the Allies and successfully execute the Blitzkrieg phase of PENETRATION and selective EXPLOITATION by MANEUVER. You can't slow an advance to defend your VULNERABLE flank by CONCENTRATION and still maintain the MOMENTUM and AVOIDANCE that come with MANEUVER.

Balance

The high pace of MANEUVERS executed by the Germans had thrown the Allies off balance. They would first turn one way in response to the Germans, who would by then have turned another way, throwing the Allies further off balance. When you are moving one way it is hard to suddenly switch directions; the Germans were choosing where to go, but the Allies had to react to the Germans out of a position of lack of knowledge (SECURITY), creating situations where the new German MOVEMENT would require a response that was harder to make just because of the one just made, IMBALANCING the Allies. The continuous imbalancing had paralyzed the Allies (DISLOCATION), leaving them unable to relationally MANEUVER to counter the Germans. By CONCENTRATING, and thus giving a more stable target to respond to, the Germans had allowed the Allies to regain their BALANCE.

Adaptation

The Allies needed to halt the German MOMENTUM. But the Germans had provided no opportunities to do this until they TIMIDLY slowed their advance in Scene-7. The Allies capitalized on the SITUATION immediately by RAPIDLY counterattacking the temporarily VULNERABLE Germans. If they had waited, this window of opportunity would have vanished as the Germans CONCENTRATED, ending their VULNERABILITY. The timely counterattack was a very EXPEDIENT action, carried out EFFICIENTLY.

Sometimes it would not be wise to attack a stronger force; in this case, the Allies needed to throw the Germans off BALANCE to get time to recover from their DISLOCATION--so the counterattack was an appropriate ADAPTATION to this particular situation.

How are SIA Themes interrelated to avoid compartmentalization?

We have shown how much is learned about the full structure of SIA in a single battle scene just by successively reading a few Core Theme commentaries for several scenes in succession, and several additional commentaries for one scene. Now we *return* to in-depth exploration of Scene 8, with the intention of lessening any *compartmentalization* that may have resulted from the list-like treatments (reading one Core Theme after another) in the preceding sections. A central tenet of Cognitive Flexibility Theory is the avoidance of compartmentalization: in ill-structured domains, knowledge is highly *interconnected* in its contexts of application. The example of the Cognitive Flexibility Hypertext for SIA presented thus far has illustrated *several anti-compartmentalization features*: Multiple conceptual themes are presented in the context of the *same case-scene*, rather than the conventional textbook treatment of separating the conceptual elements and presenting an example of each one separately. Theme commentaries have built-in *cross-references* to themes that are related in their application to the scene. Now it is time to pay more attention to the interrelatedness of the themes in a scene. We would like you to re-read the theme commentaries presented for Scene 8 (the Core Theme commentaries presented earlier, and the additional theme commentaries just presented). Except, this time, *pay particular attention to the relationships* among themes. Use any clues you notice to explore around on your own. You should notice your understanding of SIA, as it applies to Scene 8, becoming more organically interconnected, as a representational *weave* is beginning to be formed. The longer you do it, the more interconnections you will notice.

Building partial schemas and tentative generalizations: Search for cases illustrating theme combinations

In this section we will begin to see how interconnectedness gets built *across* minicase-scenes (helped by the interconnectedness of themes observed *within* scenes). One pattern of relationship that you are likely to have noticed in the in-depth exploration of Scene 8 involved the following elements:

The Germans relinquished their MOMENTUM because of TIMIDITY, reducing the DISLOCATION of the Allies and thereby enabling a LEVERAGED counter-attack (i.e., a move that had greater effects than it ought to have, given its relative lack of strength).

Is there a *general lesson* to be learned here? Is this a pattern that applies to other cases as well? One feature of Cognitive Flexibility Hypertexts is that they permit a search for any minicase-scenes that have relevance to the intersection of any *combination* of conceptual themes (a further reduction of conceptual and case *compartmentalization*, as recommended by Cognitive Flexibility Theory). This is done simply by highlighting the desired themes, the four themes capitalized above, in the Theme Menu (see Figure 40 on the next page). The search yields a scene from the Civil War battle of Chancellorsville that shows a battle that had a similar combination of SIA themes. The battle map and background and description for that scene are presented on the second, third, and fourth pages following, Figures 41-43.

Click on the themes you want to search for. Then, click Initiate Search.

Core SIA Themes	Bridge Themes	
Dislocation	Deception	Balance
Flexibility	Mystification	Disorder
Situatedness/Expediency	Distraction	Morale
Indirection	Adaptation	Boldness/Timidity
Efficiency	Momentum	Expectation
	Leverage	Multiplicity and Change
Initiate Search	Avoidance	Exploitation/Vulnerability
Return to Main Menu	Maneuver/Mobility	Dogma-boundedness
	Obliqueness	
Case Themes	Case Features	
Concentration/Dispersion	Feints, Ruses, Traps	Central axis
Commitment	Shapes opponent	Breakthrough
Speed/Expeditiousness	Zig-zag	Penetration
Movement/Stasis	Counterattack	Wedging
Initiative	Threat: frontal	Rolling out
Security	Threat: flank	Force strength
Surprise	Threat: rear & communications	Terrain
Alternative objectives	Threat: envelopment	Weather

Figure 40. Combinatorial theme search.

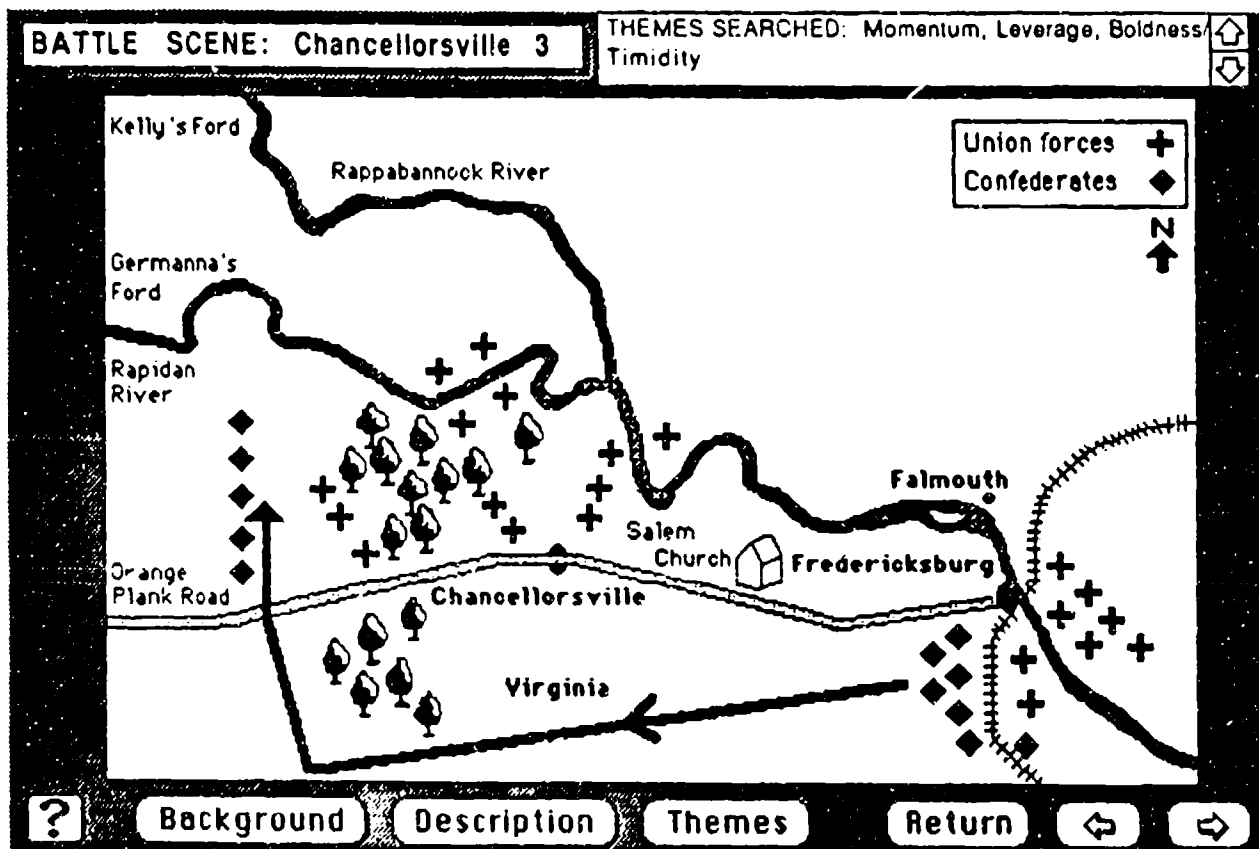


Figure 41. Chancellorsville scene

BATTLE SCENE: Chancellorsville 3

THEMES SEARCHED: Momentum, Leverage, Boldness, Timidity

BATTLE BACKGROUND

Hooker has just executed a daring flanking maneuver around the Confederate left. With his overwhelming force superiority and his strong position, he is now in an excellent position to pursue the initiative.

BATTLE DESCRIPTION

After a brilliant flanking maneuver has put him in an excellent offensive position, he does nothing. He simply remains under the cover of the Wilderness. In the meantime, Lee takes advantage of the Union Army's inactivity to boldly move across their front. He sends Jackson through the cover of the Wilderness around Hooker's unprotected right flank. Now

The map shows the battle area with Union forces marked by '+' and Confederate forces by '♦'. A dashed line indicates a path or front line. Labels 'Falmouth' and 'Federicksburg' are present. A north arrow points upwards.

?

Background

Description

Themes

Return

↶

↷

Figure 42. Chancellorsville scene: Background and description (part 1)

BATTLE SCENE: Chancellorsville 3

THEMES SEARCHED: Momentum, Leverage, Boldness, Timidity

BATTLE BACKGROUND


Hooker has just executed a daring flanking maneuver around the Confederate left. With his overwhelming force superiority and his strong position, he is now in an excellent position to pursue the initiative.

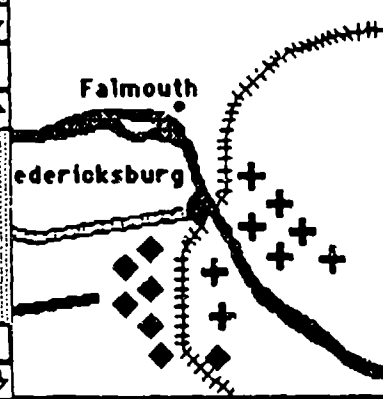
BATTLE DESCRIPTION

nothing. He simply remains under the cover of the Wilderness. In the meantime, Lee takes advantage of the Union Army's inactivity to boldly move across their front. He sends Jackson through the cover of the Wilderness around Hooker's unprotected right flank. Now the Confederate Army is in the stronger offensive position.

Union forces
+

Confederates
◆





?
Background
Description
Themes
Return
↶
↷

Figure 43. Chancellorsville scene: Background and description (part 2)

So, we now have preliminary evidence that this particular *combination* of themes might form some kind of general rule or generic schema, useful across a range of cases. (Even ill-structured domains have some well-structured elements.) To determine if that is true, the theme-combination search would be continued. If counter-examples are discovered, they could be compared with the scenes where the candidate schema works to determine whether there are other theme values (besides the four that are being tracked) which would help the learner to determine the *contextual conditions* for which the schema works and those for which it does not.

4.6.3 Concluding Remarks

At this point we will end our tour of the initial stages of learning SIA using the Cognitive Flexibility Hypertext. Of course, we have barely scratched the surface, just beginning to cycle out of the Ardennes case to consider other historical battles. Yet, much has already been learned, and you can probably imagine with some ease how the instructional process for teaching conceptual aspects of SIA in case contexts would continue. There are many other issues that we have not been able to touch on (e.g., algorithms for sequencing; options for the learner to customize the program as he or she makes new discoveries; and so on -- the unique advantages of our hypertext design approach compared to others were covered in detail in Chapter 4.4). However, our purpose was to provide a detailed concrete example to illustrate the hypertext theory that was presented in great depth in Chapter 4.4, and that has been accomplished.

Even with the amount of material that we have been able to cover, with just a partial picture of the experience of a Cognitive Flexibility Hypertext, it is clear that many Cognitive Flexibility Theory principles, necessary for advanced knowledge acquisition because of the features of ill-structuredness, have already begun to be well addressed. For example, you have been able to see how multiple themes apply to any case for fullness of case coverage and to create multiple knowledge assembly routes for future use. How conceptual variability is taught by the theme search process (and flexible future use of each theme is thereby engendered). How the themes and scenes are interrelated in various ways rather than compartmentalized. How themes are combined to produce larger conceptual structures. How the meaning of each theme gets tailored to the specific case context in which it is being applied (thereby preparing for future situated applications). And much more. (Again, see Chapter 4.4 for a more detailed accounting of the many ways that Cognitive Flexibility Hypertexts provide the features that are needed to prepare an individual to flexibly deal with a wide range of disparate cases in the future.)

Most importantly, it should be clear how, by compacting and elaborating cases without removing their *structural* features of complexity, the goal of consolidating many times over the process of knowledge acquisition from case experience is addressed (or at least that part of experience that results in *conceptual knowledge structure building*). Think of how much was taught in so short an amount of time -- and, again, we have barely scratched the surface. Imagine encountering the same kind of thing as you saw for Scene 8 for another scene, then another, then another, with each of the multi-lessoned new scenes interacting with each of those that had proceeded, and so on. Imagine what would happen as knowledge webs get tractably and incrementally built up over hours, or days, or weeks (just by adding cases; there is no "scaling-

up problem" with Cognitive Flexibility Hypertexts--see below). Cognitive Flexibility Hypertexts are supposed to be used over periods of weeks, not an hour or two. And, of course, advanced case-based instruction over days or weeks is far more desirable than the many years it usually takes to acquire sufficient everyday experience for the attainment of expertise in ill-structured domains. The effectiveness of the SIA hypertext has not been tested yet. Programmatically, it made the most sense to begin with domain development and tests using more manageable domains and learning materials. However, the preliminary indications of transfer success when following Cognitive Flexibility Theory principles with the simpler materials, and the considerable overlap between the earlier domains and SIA in the way Cognitive Flexibility Theory is applied, provide good grounds for optimism.

4.6.3.1 A Final Note on the Difficulty of Preparing Texts for Cognitive Flexibility Hypertexts

It may seem at first that all of the text writing you have observed would make it difficult to build a Cognitive Flexibility Hypertext. This is not a very big problem for several reasons. *First*, other cases are not covered in as much detail as Ardennes, making them easier to produce. Ardennes is intended to be the *starting point* of instruction in SIA. This follows the Cognitive Flexibility Theory principle of the New Incrementalism: start case explorations with "bite-size chunks of complexity" (here, a great deal of weaving of very compact scenes). *Second*, what was written just for the Ardennes case bought a large amount of potential lessons and learning time. Even after you have used the materials you read for an hour or two, as described, you could still return to them many more times to explore them in different criss-crossed sequences, each teaching new lessons about the complex SIA. And the "fanciness" of what Cognitive Flexibility Hypertexts do lies more in the number of weaves within and between many minicases, along many conceptual dimensions, not in the complexity or subtlety of the individual scene commentaries. *Third*, each new case, even if not as detailed as Ardennes, adds multiple minicase scenes and an exponential explosion of potential interconnections when a highly webbed, multi-thematic new case and its scenes gets woven through an increasingly more substantial preexisting web of case-based SIA knowledge. *Fourth*, each scene or case *ends* a "story", so you can start anywhere -- i.e., you have completeness or closure after writing just one case. And you can add cases *any time*, gradually building up a large set of cases--the more experience the better in an ill-structured domain--while being able to do plenty with the cases you already have at every step of the way). *Fifth*, there are no *scaling-up* problems that always plague computer instructional systems as they move from initial prototypes -- all you do to expand the program is *more of the same*--everything has been done that's difficult to do; and all learners have to do is the same thing they have already done for new cases, always increasing the number of revisitations of previously read material (although decreasing the prob. of revisit per case/scene). Again, other virtues of the Cognitive Flexibility Hypertext design approach and factors contributing to its efficiency and practicability are dealt with in detail in Chapter 4.4.

5. Acknowledgments

The authors would like to acknowledge the contributions of the following individuals to the conduct of segments of this research: Rita Gaskill, Jane Adami, Joan Feltovich, and Kathy Cain.

Most of all, we wish to express our gratitude to the professional staff of the Office of Basic Research of the U.S. Army Research Institute for the Behavioral and Social Sciences. Michael Kaplan (Director), Judith Orasanu, Michael Drillings, and George Lawton have helped in innumerable ways to facilitate this research. Their frequent input and advice have been immeasurably helpful in shaping the directions of this work and, I believe very strongly, in improving it. One could not ask to collaborate with a more thoughtful, forward-looking, and understanding group of individuals. Of course, the contents of this report represent the thoughts solely of the authors and do not reflect the views of the agency.

A note on the report: In various places in this report, material from papers authored by the Principal Investigator with the support of this ARI grant is incorporated directly. Thus large parts of the material in the following chapters is drawn from the respective published paper supported by this grant: 2.0: Spiro et al., 1988; 3.1: Feltovich et al., 1969; 3.2: Spiro et al., 1989; 4.1: Spiro et al., 1987; 4.2: Spiro et al., 1989; 4.3: Spiro et al., in press; 4.4: Spiro & Jehng, 1990; 4.5: Jacobson & Spiro, in press. The individuals who were primary contributors to each chapter in the report are listed in a footnote following the chapter. Also, elements of some of the work discussed in this report benefited from other sources of support, in addition to ARI, and are included either for reasons of continuity of argument or because the flow-through of developments from one program of study to another produced a pattern of mutual benefit. This is especially so for the Office of Educational Research and Improvement role in parts of 4.1 and 4.5, and the Macy Foundation role in parts of 3.1. Agencies that contributed to smaller parts of other work reported herein, complementing the funding of ARI, include: Office of Educational Research and Improvement, Office of Naval Research, and International Business Machines Corp. A Department of Defense "DURIP" equipment grant to the P.I., related to the ARI grant, was another major source of support.

6. Papers and Presentations of Research Supported by the Grant

Publications

- Spiro, R. J., Coulson, R. L., Feltovich, P. J., & Anderson, D. (1988). Cognitive flexibility theory: Advanced knowledge acquisition in ill-structured domains. *Tenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Coulson, R. L., Feltovich, P. J., & Spiro, R. J. (1989). Foundations of a misunderstanding of the ultrastructural basis of myocardial failure: A reciprocating network of oversimplifications. *Journal of Medicine and Philosophy*, 14, 109-146. [Special issue on "The Structure of Clinical Knowledge."]
- Feltovich, P. J., Spiro, R. J., & Coulson, R. L. (1989). The nature of conceptual understanding in biomedicine: The deep structure of complex ideas and development of misconceptions. In D. Evans & V. Patel (Eds.), *The cognitive sciences in medicine* (pp. 113-172). Cambridge, MA: MIT Press.
- Spiro, R. J., Feltovich, P. J., Coulson, R. L., & Anderson, D. (1989). Multiple analogies for complex concepts: Antidotes for analogy-induced misconception in advanced knowledge acquisition. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 498-531). Cambridge: Cambridge University Press.
- Feltovich, P. J., Spiro, R. J., & Coulson, R. L. (in press). Learning, teaching and testing for complex conceptual understanding. In N. Frederiksen, R. Mislevy, & I. Bejar (Eds.), *Test theory for a new generation of tests*. Hillsdale, NJ: Erlbaum.
- Spiro, R. J., Feltovich, P. J., Jacobson, M. J., & Coulson, R. L. (in press). Cognitive flexibility, constructivism, and hypertext: Random access instruction for advanced knowledge acquisition in ill-structured domains. *Educational Technology*.
- Jacobson, M. J., & Spiro, R. J. (in press). Hypertext learning environments, cognitive flexibility, and the transfer of complex knowledge: An empirical investigation. *Proceedings of the Conference of the Learning Sciences*.
- Spiro, R. J., & Jehng, J. C. (1990). Cognitive flexibility and hypertext: Theory and technology for the nonlinear and multidimensional traversal of complex subject matter. In D. Nix, & R. J. Spiro (Eds.), *Cognition, education, and multimedia: Exploring ideas in high technology* (pp. 163-205). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Nix, D., & Spiro, R. J. (Eds.) (1990). *Cognition, education, and multimedia: Explorations in high technology*. Hillsdale, NJ: Lawrence Erlbaum.

Talks

Spiro, R. J., Jacobson, M., & Jehng, J. (November 1988). *Hypertext and cognitive psychology: Theory and technology for computer-based instruction in complex knowledge domains*. Paper presented at the Annual Meeting of the Association for the Development of Computer-Based Instructional Systems (ADCIS), Philadelphia.

Spiro, R. J., Jehng, J., Durgunoglu, A., McGinley, W., & Jacobson, M. (December 1988). *Beyond schema theory: Developing flexible background knowledge for independent reading comprehension in the content areas*. Invited symposium paper presented at the Annual Meeting of the National Reading Conference, Tucson.

Hartman, D. K., & Spiro, R. J. (March 1989). *Explicit text structure instruction for advanced knowledge acquisition in complex domains: A poststructuralist perspective*. Paper presented at the Annual Meeting of the American Educational Research Association, San Francisco.

Spiro, R. J. (March 1989). *Professional cognition: The management of complexity*. Division-I invited address presented at the Annual Meeting of the American Educational Research Association, San Francisco.

Spiro, R. J. (April 1989). *Learning and instruction in complex knowledge domains*. Colloquium presented at the Pennsylvania State University.

Spiro, R. J. (April 1989). *Knowledge acquisition in complex domains: Hypertext and the development of cognitive flexibility*. Plenary address at the Ninth Conference on Reading Research (CORR9), International Reading Association, New Orleans.

Spiro, R. J. (June 1989). *Advanced skills training*. Colloquium presented at the U.S. Army Training and Doctrine Command, Fort Monroe, Virginia.

Spiro, R. J. (1990, March). *Knowledge acquisition in complex domains: Hypertext and the development of cognitive flexibility*. Colloquium presented at the University of Michigan.

Spiro, R. J. (1990, April). *Toward a general theory of hypertext: A nonlinear instruction and the development of cognitive flexibility in ill-structured domains*. Invited paper presented at NATO Advance Research Workshop on "Structure of Communication and Intentional Help for Hypertext Courseware," Espinho, Portugal.

Spiro, R. J. (1990, April). *Learning in complex domains*. Fulbright Lectures presented at the University of Aveiro and the University of Coimbra, Portugal.

Spiro, R. J., (1990, May). *The acquisition and application of knowledge in complex domains*. Colloquium presented at the University of Georgia.

- Spiro, R. J. (1990, August). *Learning in ill-structured domains: Cognitive flexibility and the management of complexity*. Invited address presented at the annual meeting of the American Psychology Association, Boston. [Division 3 (Experimental Psychology) and 15 (Educational Psychology)]
- Spiro, R. J., Feltovich, P. J., Jacobson, M. J., & Coulson, R. L. (in press). Cognitive flexibility, constructivism, and hypertext: Random access instruction for advanced knowledge acquisition in ill-structured domains. To appear in *Educational Technology* (Special Issue on Constructivism).
- Spiro, R. J., & Jehng, J. C. (1990). Cognitive flexibility and hypertext: Theory and technology for the nonlinear and multidimensional traversal of complex subject matter. In D. Nix & R. J. Spiro (Eds.), *Cognition, education, and multimedia: Explorations in high technology* (pp. 163-205). Hillsdale, NJ: Lawrence Erlbaum.

7. References

- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Anderson, J. R., Kline, P. G., & Beasley, C. M. (1979). Schemata as scaffolding for the representation of information in discourse. *American Educational Research Journal*, 15, 433-440.
- Anderson, R. C. (1977). The notion of schema and the educational enterprise. In R. C. Anderson, R. J. Spiro, & W. E. Montague (Eds.), *Schooling and the acquisition of knowledge*. Hillsdale, NJ: Erlbaum.
- Anderson, R. C., Spiro, R. J., & Anderson, M. C. (1978). Schemata as scaffolding for the representation of information in discourse. *American Educational Research Journal*, 15, 433-440.
- Associated Press dispatch (dispatch reported in Champaign-Urbana News-Gazette), 1990. [Exact date not known; a copy of the article is available.]
- Ausubel, D. P. (1968). *Educational psychology: A cognitive view*. New York: Holt, Rinehart and Winston.
- Barrows, H. S., & Tamblyn, R. M. (1980). *Problem-based learning*. New York: Springer.
- Barsalou, L. W. (1987). Instability of graded structure: Implications for the nature of concepts. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 101-140). Cambridge: Cambridge University Press.
- Barthes, R. (1967). *S/Z*. New York: Hill, Wang.
- Bartlett, F.C. (1932). *Remembering*. Cambridge, UK: Cambridge Univ. Press.
- Bednar, A. K., Cunningham, D., Duffy, T.M., Perry, J.D. (1991). Theory into practice: How do we link? In G. Anglin (Ed.), *Instructional technology: Past, present, and future*. Denver, CO: Libraries Unlimited.
- Beeman, W. O., Anderson, K. T., Bader, G., Larkin, J., McClard, A. P., McQuillan, P. J., & Shields, M. (1987). Hypertext and pluralism: From lineal to non-lineal thinking. In *Hypertext '87 Papers* (pp. 67-88). Chapel Hill: The University of North Carolina
- Beeman, W. O., Anderson, K., Bader, G., Larkin, J., McClard, A. P., McQuillan, P., & Shields, M. (1988). *Intermedia: A case study of innovation in higher education* (Final report to the Annenberg/CPB Project). Providence, RI: Brown University, Office of Program Analysis, Institute for Research in Information and Scholarship.

- Bereiter, C. (1985). Toward a solution of the learning paradox. *Review of Educational Research*, 55, 201-226.
- Bereiter, C., & Scardamalia, M. (1985). Cognitive coping strategies and the problem of "inert knowledge." In S. F. Chipman, J. W. Segal, & R. Glaser (Eds.), *Thinking and learning skills: Current research and open questions*. (Vol. 2, pp. 65-80). Hillsdale, NJ: Lawrence Erlbaum.
- Bransford, J. D., Franks, J. J., Vye, N. J., & Sherwood, R. D. (1989). New approaches to instruction: Because wisdom can't be told. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (470-497). Cambridge, MA: Harvard University Press.
- Bransford, J. D., & Johnson, M. K. (1972). Contextual prerequisites for understanding: Some investigations of comprehension and recall. *Journal of Verbal Learning and Verbal Behavior*, 11, 717-726.
- Bransford, J. D., Nitsch, K. E., & Franks, J. J. (1977). Schooling and the facilitation of knowing. In R. C. Anderson, R. J. Spiro & W. E. Montague (Eds.), *Schooling and the acquisition of knowledge*. Hillsdale, NJ: Erlbaum.
- Britton, B. K., & Black, J. B. (1985). *Understanding expository text*. Hillsdale, NJ: Erlbaum.
- Brown, A. L. (1989). Analogical learning and transfer: What develops? In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (369-412). Cambridge, MA: Harvard University Press.
- Brown, A.L. & Campione, J.C. (1981). Inducing flexible thinking: A problem of access. In M. Friedman, J.P. Das, & N. O'Connor (Eds.), *Intelligence and learning* (515-529). New York: Plenum.
- Brown, A. L., Campione, J. C., & Day, J. (1981). Learning to learn: On training students to learn from texts. *Educational Researcher*, 10, 14-21.
- Brown, J.S., Collins, A., & Duguid, P. (1989). Situated cognition and the culture of learning. *Educational Researcher*, 18, 32-42.
- Bruner, J.S. (1963). *The process of education*. Cambridge, MA: Harvard University Press.
- Burstein, M. (1983). *Concept formation by incremental analogical reasoning and debugging*. Proceedings of the International Machine Learning Workshop, Champaign, IL.
- Burstein, M.H., & Adelson, B. (1990). Issues for a theory of analogical learning. In R. Freedle (Ed.), *Artificial intelligence and the future of testing*. Hillsdale, NJ: Lawrence Erlbaum Associates.

- Carbonell, J. G. (1983). *Derivational analogy in problem solving and knowledge acquisition*. Proceedings of the International Machine Learning Workshop, Champaign, IL.
- Carnine, D. (1980). Three procedures for presenting minimally different positive and negative instances. *Journal of Educational Psychology*, 7, 452-456.
- Burstein, M. (1983). *Concept formation by incremental analogical reasoning and debugging*. Proceedings of the International Machine Learning Workshop, Champaign, IL.
- Champagne, A. B., Gunstone, R. F., & Klopfer, L. E. (1985). Effecting changes in cognitive structures among physics students. In L. H. West & A. L. Pines (Eds.), *Cognitive structure and conceptual change*. Orlando, FL: Academic Press.
- Chase, W.C., & Simon H. (1973). Perception in chess. *Cognitive Psychology*, 4, 55-81.
- Chernoff, I. (1973). The use of faces to represent points in k-dimensional space graphically. *Journal of the American Statistical Association*, 68, 361-368.
- Chi, M.T.H., Feltovich, P.J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 5, 121-152.
- Chi, M. T. H., Glaser, R., & Farr, M. J. (Eds.). (1988). *The nature of expertise*. Hillsdale, NJ: Lawrence Erlbaum.
- Collins, A., Brown, J. S., & Newman, S. E. (in press). Cognitive apprenticeship: Teaching the craft of reading, writing, and mathematics. In L. B. Resnick (Ed.), *Cognition and instruction: Issues and agenda*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Collins, A., & Gentner, D. (in press). How people construct mental models. In S. Quinn & D. Holland (Eds.), *Cultural models in thought and language*. Cambridge: Cambridge University Press.
- Collins, A., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82, 407-428.
- Collins, A., & Stevens, A. L. (1983). Goals and strategies of inquiry teachers. In R. Glaser (Ed.), *Advances in instructional psychology* (Vol. 2). Hillsdale, NJ: Erlbaum.
- Conklin, J. (1987). Hypertext: An introduction and survey. *IEEE Computer*, 20(9), 17-41.
- Coulson, R. L., Feltovich, P. J., & Spiro, R. J. (1987). *A test of biomedical reasoning and conceptual understanding based on a theory of reasoning about complex systems*. Manuscript in preparation.

- Coulson, R.L., Feltovich, P.J., & Spiro, R.J. (1989). Foundations of a misunderstanding of the ultrastructural basis of myocardial failure: A reciprocation network of oversimplifications. *The Journal of Medicine and Philosophy (special issue on "The Structure of Clinical Knowledge")*, 14, 109-146.
- Crane, G., & Mylonas, E. (1988). The Perseus project: Interactive curriculum on classical Greek civilization. *Educational Technology*, 28(11), 25-32.
- Dawson-Saunders, B. K., Feltovich, P. J., Coulson, R. L., & Steward, D. (1987). *The current "top 20" hits: Biomedical concepts all medical students need to understand*. Paper presented at the Annual Meeting of the American Educational Research Association, Washington, D.C.
- Dede, C. J. (1987). Empowering environments, hypermedia and microworlds. *The Computing Teacher*, 15(3), 20-24, 61.
- Dede, C. J. (1988). The probable evolution of artificial intelligence based educational devices. *Technological Forecasting and Social Change*, 34, 115-133.
- Elstein, A. S. (in press). Cognitive processes in clinical inference and decision making. In D. Turk & P. Salovey (Eds.), *Reasoning, inference, and judgment in clinical psychology*. New York, NY: Free Press/Macmillan.
- Feltovich, P. J., Coulson, R. L., & Spiro, R. J. (1986). *The nature and acquisition of faulty student models of selected medical concepts: Cardiovascular impedance*. Paper presented at the Annual Meeting of the American Educational Research Association, Montreal.
- Feltovich, P. J., Coulson, R. L., & Spiro, R. J. (in press). The nature of conceptual understanding in biomedicine: The deep structure of complex ideas and the development of misconceptions. In D. Evans & V. Patel (Eds.), *The cognitive sciences in medicine*. Cambridge, MA: MIT Press.
- Feltovich, P.J., Johnson, P.E., Moller, J.H., & Swanson, D.B. (1984). LCS: The role and development of medical knowledge in diagnostic expertise. In W. J. Clancey & E. H. Shortliffe (Eds.), *Readings in medical artificial intelligence: The first decade*. Reading, MA: Addison Wesley.
- Feltovich, P.J., Spiro, R.J., & Coulson, R.L., (1989). The nature of conceptual understanding in biomedicine: The deep structure of complex ideas and the development of misconceptions. In D. Evans & V. Patel (Eds.), *Cognitive science in medicine* (pp. 113-172). Cambridge, MA: MIT Press (Bradford Books).
- Feltovich, P. J., Spiro, R. J., & Coulson, R. L. (in press). Learning, teaching, and testing for complex conceptual understanding. In N. Frederiksen, R. Missevsky, & I. Bejar (Eds.), *Test theory for a new generation of tests*. Hillsdale, NJ: Lawrence Erlbaum Associates.

- [GPEP] Panel on the General Professional Education of the Physician. (1984). *Physicians for the twenty-first century: The GPEP report*. Washington, DC: American Association of Medical Colleges.
- Gentner, D. (1983). Structure mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.
- Gentner, D., & Gentner, D. R. (1983). Flowing waters or teeming crowds: Mental models of electricity. In D. Gentner & A. L. Stevens (Eds.), *Mental models*. Hillsdale, NJ: Erlbaum.
- Gibbs, C., & Loiselle, D. (1978). The energy output of tetanized cardiac muscle: Species differences. *Pflügers Arch*, 373, 31-38.
- Gick, M. L., & Holyoak, K. J. (1987). The cognitive basis of knowledge transfer. In S. M. Cormier, and J. D. Hagman (Eds.), *Transfer of learning: Contemporary research and applications* (pp. 9-46). New York : Academic Press.
- Gick, M. L., & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology*, 12, 306-355.
- Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, 14, 1-38.
- Glaser, R. (1984). Education and thinking: The role of knowledge. *American Psychologist*, 39, 93-104.
- Gleick, J. (1987). Snowflake study unlocks pattern formation mysteries. *The State Journal Register* (Springfield, IL), 25.
- Goldberger, E. (1982). *Textbook of clinical cardiology*. St. Louis, MO: C. V. Mosby.
- Gordon, A. M., Huxley, A. F., & Julian, F. J. (1966). The variation in isometric tension with sarcomere length in vertebrate muscle fibers. *Journal of Physiology (Lond)*, 184, 170-192.
- Hart, B. H. L. (1954). *Strategy*. London: Faber & Faber.
- Hartman, D. K., & Spiro, R. J. (1989). *Explicit text structure instruction for advanced knowledge acquisition in complex domains: A post-structuralist perspective*. Paper presented at the annual meeting of the American Educational Research Association, San Francisco, CA.
- Huntsman, L. L., Rondinone, J. F., & Martyn, D. A. (1983). Force length relations in cardiac muscle segments. *American Journal of Physiology*, 244, H701-H707.

- Jacobson, M. J., & Spiro, R. J. (1991a). *Cognitive flexibility, hypertext learning environments, and the acquisition of complex knowledge: An empirical investigation*. Manuscript submitted for publication.
- Jacobson, M. J., & Spiro, R. J. (1991b). *A framework for the contextual analysis of computer-enhanced learning environments*. Manuscript submitted for publication.
- Jacobson, M. J., & Spiro, R. J. (1991c). *Hypertext learning environments, cognitive flexibility, and the transfer of complex knowledge: An empirical investigation*. Manuscript submitted for publication.
- Jewell, B. R. (1977). A reexamination of the influence of muscle length on myocardial performance. *Circulation Research*, 40, 221-230.
- Jonassen, D. H. (1986). Hypertext principles for text and courseware design. *Educational Psychologist*, 21(4), 269-292.
- Jonassen, D. H. (1988). Designing structured hypertext and structuring access to hypertext. *Educational Technology*, 28(11), 13-16.
- Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgment under uncertainty*. New York: Cambridge University Press.
- Katz, A. M. (1977). *Physiology of the heart*. New York, NY: Raven Press.
- Kolodner, J. L. (1980). *Retrieval and organizational strategies in conceptual memory: A computer model* (Research Rep. 187). New Haven, CT: Department of computer Science, Yale University.
- Kuhn, T. S. (1971). *The structure of scientific revolutions* (3rd edition). Chicago, IL: University of Chicago Press.
- Larkin, J. H., McDermott, J., Simon, D. P., & Simon, H. A. (1980). Expert and novice performance in solving physics problems. *Science*, 208, 1335-1342.
- Lenat, D. B., Hayes-Roth, F., & Klahr, P. (1983). *Cognitive economy in fluid task environment*. Proceedings of the International Machine Learning Workshop, Champaign, IL.
- Luttwak, E. N. The operational level of war. In S. E. Miller (Ed.), *Conventional forces and American defense policy* (pp.211-230). Princeton, N.J.: Princeton Univ. Press.
- Marchionini, G. (Ed.). (1988). Special issue on hypermedia [Special issue]. *Educational Technology*, 29(11).

- McClelland, J. L., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, 114, 159-188.
- McClelland, J. L., Rumelhart, D. E., & Hinton, G. E. (1986). The appeal of parallel distributed processing. In D. E. Rumelhart, J. L. McClelland & the PDP Research Group (Eds.), *Parallel distributed processing (Vol. 1: Foundations)*. Cambridge, MA: MIT (Bradford) Press.
- McClelland, J. L., Rumelhart, D. E., & the PDP Research Group (1986). *Parallel distributed processing (Vol. 2: Psychological and biological models)*. Cambridge, MA: MIT (Bradford) Press.
- McCloskey, M. (1983). Naive theories of motion. In D. Gentner and A. L. Stevens (Eds.), *Mental models*. Hillsdale, NJ: Lawrence Erlbaum.
- Medin, D. L., & Schaffer, M. H. (1978). A context theory of classification learning. *Psychological Review*, 85, 207-238.
- Medin, D. L., & Schwanenflugel, P. J. (1981). Linear separability in classification learning. *Journal of Experimental Psychology: Human Learning and Memory*, 7, 355-368.
- Medin, D. L., & Smith, E. E. (1984). Concepts and concept formation. *Annual Review of Psychology*, 35, 113-138.
- Merrill, M. D. (1983). Component display theory. In C. M. Reigeluth (Ed.), *Instructional-design theories and models: An overview of their current status*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Minsky, M. (1975). A framework for representing knowledge. In P. Winston (Ed.), *The psychology of computer vision*. New York, NY: McGraw-Hill.
- Minsky, M. (1979). The society theory. In P. Winston (Ed.), *Artificial Intelligence: An MIT perspective*. Cambridge, MA: MIT Press.
- Myers, A.C., Feltoovich, P.J., Coulson, R.L., Adami, J.F., & Spiro, R.J. (1990). Reductive biases in the reasoning of medical students: An investigation in the domain of acid-base balance. In B. Bender, R.J. Hiemstra, A.J.J.A. Scherbier, & R.P. Zwierstra (Eds.), *Teaching and assessing clinical competence*. Groningen, the Netherlands: BoekWerk Publications.
- Nair, K. G., Cutliff, A. F., Zal, R., et al. (1968). Biochemical correlates of cardiac hypertrophy. 1. experimental model: Changes in heart weight, RNA content and nuclear RNA polymerase activity. *Circulation Research*, 23, 451-462.
- Pepper, S. (1942). *World hypotheses*. Berkeley: University of California Press.

- Perkins, D.N., & Simmons, R. (1989). Patterns of misunderstanding: An integrative model for science, math, and programming. *Review of Educational Research*, 58, 303-326.
- Rosch, E. H., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7, 573-605.
- Ross, B. H. (1984). Reminders and their effects in learning a cognitive skill. *Cognitive Psychology*, 16, 371-416.
- Rumelhart, D. E. (1980). Schema: The building block of cognition. In R. J. Spiro, B. C. Bruce, & W. F. Brewer (Eds.), *Theoretical issues in reading comprehension*. Hillsdale, NJ: Erlbaum.
- Rumelhart, D. E., McClelland, J. L., & the PDP Research Group (Eds.). (1986). *Parallel distributed processing (Vol. 1: Foundations)*. Cambridge, MA: MIT (Bradford) Press.
- Rumelhart, D. E., Hinton, G. E., & McClelland, J. L. (1986). A general framework for parallel distributed processing. In D. E. Rumelhart, J. L. McClelland and the PDP Research Group (Eds.), *Parallel distributed processing (Vol. 1: Foundations)*. Cambridge, MA: MIT (Bradford) Press.
- Rumelhart, D. E., & Norman, D. A. (1981). Analogical processes in learning. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition*. Hillsdale, NJ: Erlbaum.
- Rumelhart, D. E., & Ortony, A. (1977) The representation of knowledge in memory. In R. C. Anderson, R. J. Spiro, & W. E. Montague (Eds.), *Schooling and the acquisition of knowledge*. Hillsdale, NJ: Erlbaum.
- Schank, R. C. (1982). *Dynamic memory*. Cambridge: Cambridge University Press.
- Schank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals, and understanding*. Hillsdale, NJ: Erlbaum.
- Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review*, 84, 1-66.
- Schoenfeld, A. H. (1983). Beyond the purely cognitive: Belief systems, social cognitions, and metacognitions as driving forces in intellectual performance. *Cognitive Science*, 7, 329-363.
- Scott, W. A., Osgood, D. W., & Peterson, C. (1979). *Cognitive structure: Theory and measurement of individual differences*. Washington, DC: Winston.
- Shepard R. N. (1967). Recognition memory for words, sentences, and pictures. *Journal of Verbal Learning and Verbal Memory*, 6, 156-163.

- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, 84, 127-190.
- Siegler, R. S., & Klahr, D. (1982). When do children learn? The relationship between existing knowledge and the acquisition of new knowledge. In R. Glaser (Eds.), *Advances in instructional psychology* (Vol. 1). Hillsdale, NJ: Erlbaum.
- Smith, E. E., & Medin, D. L. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.
- Smith, S., Glenberg, A., & Bjork, R. A. (1978). Environmental context and human memory. *Memory and Cognition*, 6, 342-353.
- Smolensky, P. (1986). Information processing in dynamical systems: Foundations of harmony theory. In J. L. McClelland, D. E. Rumelhart & the PDP Research group (Eds.), *Parallel distributed processing* (Vol. 2). Cambridge, MA: MIT (Bradford) Press.
- Spiro, R. J. (1977). Remembering information from text: The "state of the schema" approach. In R. C. Anderson, R. J. Spiro & W. E. Montague (Eds.), *Schooling and the acquisition of knowledge*. Hillsdale, NJ: Erlbaum.
- Spiro, R. J. (1980a). Accommodative reconstruction in prose recall. *Journal of Verbal Learning and Verbal Behavior*, 19, 84-95.
- Spiro, R. J. (1980b). Constructive processes in prose comprehension and recall. In R. J. Spiro, B. C. Bruce, & W. F. Brewer (Eds.), *Theoretical issues in reading comprehension* (pp. 245-278). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Spiro, R. J. (1980). Prior knowledge and the processing of stories: Integration, selection, and variation. *Poetics: International Review for the Theory of Literature*, 9, 313-328.
- Spiro, R. J. (1983). Subjectivity and memory. In W. Kintsch & J. Le Ny (Eds.), *Language comprehension*. The Hague: North-Holland.
- Spiro, R. J. (In press). Building flexible knowledge structures from and for reading. In R. J. Spiro (Ed.), *Reading research into the 90's*. Hillsdale, NJ: Erlbaum.
- Spiro, R. J., Coulson, R. L., Feltovich, P. J., & Anderson, D. K. (1988). Cognitive flexibility theory: Advanced knowledge acquisition in ill-structured domains. In *Tenth Annual Conference of the Cognitive Science Society* (pp. 375-383). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Spiro, R. J., Crismore, A., & Turner, T. J. (1982). On the role of pervasive experiential coloration in memory. *Text*, 2, 253-262.

- Spiro, R. J., & Esposito, J. (1981). Superficial processing of explicit inferences in text. *Discourse Processes*, 4, 313-322.
- Spiro, R. J., Esposito, J., & Vondruska, R. (1978). The representation of derivable information in memory: When what might have been left unsaid is said. In D. L. Waltz (Ed.), *Theoretical issues in natural language processing (TINLAP-2)*. New York: Association for Computational Linguistics and Association for Computing Machinery.
- Spiro, R. J., Feltovich, P. J., & Coulson, R. L. (1988). *Seductive reductions: e hazards of oversimplification of complex concepts* (Tech. Rep.. 4). Springfield: Southern Illinois University Medical School.
- Spiro, R.J., Feltovich, P.J., Coulson, R.L., & Anderson, D.K. (1989). Multiple analogies for complex concepts: Antidotes for analogy-induced misconception in advanced knowledge acquisition. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (498-531). Cambridge, Mass.: Cambridge University Press.
- Spiro, R. J., Feltovich, P. J., Jacobson, M. J., & Coulson, R. L. (in press). Cognitive flexibility, constructivism, and hypertext: Random access instruction for advanced knowledge acquisition in ill-structured domains. *Educational Technology*.
- Spiro, R. J., & Jehng, J. C. (1990). Cognitive flexibility and hypertext: Theory and technology for the nonlinear and multidimensional traversal of complex subject matter. In D. Nix, & R. J. Spiro (Eds.), *Cognition, education, and multimedia: Exploring ideas in high technology* (pp. 163-205). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Spiro, R. J., & Myers, A. (1984). Individual differences and underlying cognitive processes in reading. In P. D. Pearson (Ed.), *Handbook of research in reading* (pp. 471-501). New York: Longman.
- Spiro, R. J., & Taylor, B. M. (in press). On investigating children's transition from narrative to expository discourse: The multidimensional nature of psychological text classification. In R. Tierney, J. Mitchell, & P. Anders (Eds.), *Understanding readers' understanding*. Hillsdale, NJ: Erlbaum.
- Spiro, R.J., Vispoel, W., Schmitz, J., Samarapungavan, A., & Boerger, A. (1987). Knowledge acquisition for application: Cognitive flexibility and transfer in complex content domains. In B.C. Britton & S. Glynn (Eds.), *Executive control processes in reading* (177-199). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Spotnitz, H. M., & Sonnenblick, E. H. (1976). Structural conditions in the hypertrophied and failing heart. In D. T. Mason (Ed.), *Congestive heart failure*. New York, NY: Dun-Donnelley.
- Standing, L., Conezio, J., & Haber, R. N. (1970). Perception and memory for pictures. Single-trial learning of 2500 visual stimuli. *Psychonomic Science*, 19, 73-74.

- Streufert, S., & Streufert, S. C. (1978). *Behavior in the complex environment*. Washington, DC: Winston.
- Tennyson, R.D., & Park, O. (1980). The teaching of concepts: A review of instructional design research literature. *Review of Educational Research*, 50, 55-70.
- Vosniadou, S., & Ortony, A. (1983). The influence of analogy in children's acquisition of new information from text: An exploratory study. In P. A. Niles & L. A. Harris (Eds.), *Searches for meaning in reading: Language processing and instruction*. Clemson, SC: National Reading Conference.
- Wainer, H., & Thisser, D. (1981). Graphical data analysis. *Annual Review of Psychology*, 32, 191-242.
- Waltz, D. L. (Ed.) (1985). Connectionist models and their applications [special issue]. *Cognitive Science*, 9(1).
- Waltz, D. L., & Pollack, J. (1985). Massively parallel parsing. *Cognitive Science*, 9, 51-74.
- Welles, O. (Producer/Director) & Stein, R. (Videodisc Designer and Producer) (1984). *Citizen Kane* (Videodisc No. Catalog No. CC101). Los Angeles, CA: The Criterion Collection.
- White, B. Y. (1984). Designing computer games to help physics students understand Newton's Law of Motion. *Cognition and Instruction*, 1, 69-108.
- White, B. Y., & Frederiksen, J. R. (1987, November). *Causal model progressions as a foundation for intelligent learning environments* (BBN Rep. No. 6686). Cambridge, MA: Bolt Beranek and Newman, Inc.
- White, H. (1973). *Metahistory*. Baltimore, MD: Johns-Hopkins University Press.
- Whitehead, A.N. (1929). *The aims of education*. New York: Macmillan.
- Winston, P. H. (1983). *Learning by augmenting rules and accumulating censors*. Proceedings of the International Machine Learning Workshop, Champaign, IL.
- Wittgenstein, L. (1953). *Philosophical investigations*. New York: Macmillan.
- Yates, F. A. (1966). *The art of memory*. Chicago: University of Chicago Press.

920428